

REPORT DOCUMENTATION PAGE	Form Approved OMB NO. 0704-0188
---------------------------	------------------------------------

1. AGENCY USE ONLY	2. REPORT DATE: May 13, 2001	3. REPORT TYPE: Final 01 Sep 93 - 31 Mar 98
--------------------	------------------------------	--

4. TITLE AND SUBTITLE: "F-RISC - A 1.0 GOPS FAST REDUCED INSTRUCTION SET COMPUTER FOR SUPER WORKSTATION AND TERAOPS PARALLEL PROCESSOR APPLICATIONS"	5. FUNDING NUMBERS: AMXRO-ICA #32406-EL, ARO/DARPA DAAH 04-93G-0477
6. AUTHOR(S): John F. McDonald	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES): Rensselaer Polytechnic Institute, Troy, NY 12181	8. PERFORMING ORGANIZATION REPORT NUMBER: FRISC-01
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES): U. S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	10. SPONSORING / MONITORING AGENCY REPORT NUMBER 32406-1-C1

11. SUPPLEMENTARY NOTES: The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by the documentation.


12 a. DISTRIBUTION / AVAILABILITY Approved for public release; distribution unlimited.	12 b. DISTRIBUTION CODE
--	-------------------------

13. ABSTRACT (Maximum 200 words) The purpose of this contract has been to establish whether Heterojunction Bipolar Transistors or HBT's could be used in design of high clock rate digital computers. Additionally to establish whether HBT devices offer alternatives in case conventional COTS CMOS runs into manufacturing difficulty or fundamental device limitations below 0.1 microns minimum feature size. A demonstration GaAs HBT byteallice chipset for 2 GHz 1 GOPS Fast RISC computer has been fabricated.
--

14. SUBJECT TERMS	15. # OF PAGES: 139
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT: UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE: UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT: UNCLASSIFIED	20. LIMITATION OF ABSTRACT: UL
--	---	--	-----------------------------------

NSN 7540-01-280-5500 Standard Form 298 (Rev.2-89) Prescribed by 239-18 298-102



20010608 106

**F-RISC -
“A 1.0 GOPS FAST REDUCED INSTRUCTION SET
COMPUTER FOR SUPER WORKSTATION AND
TERAOPS PARALLEL PROCESSOR APPLICATIONS”**

FINAL CONTRACT REPORT

JOHN F. MC DONALD

April 25, 2001

U. S. ARMY RESEARCH OFFICE

GRANT IDENTIFICATION NUMBERS: DAAH 04-93G-0477

[ARO #32406-EL, ARPA # A759]

RENSSELAER POLYTECHNIC INSTITUTE

**CENTER FOR INTEGRATED ELECTRONICS
AND ELECTRONICS MANUFACTURING**


TROY, NEW YORK 12180

(518) 276-2919

EMAIL: MCDONALD@UNIX.CIE.RPI.EDU

APPROVED FOR PUBLIC RELEASE

DISTRIBUTION UNLIMITED



Rensselaer

**THE VIEWS, OPINIONS, AND/OR FINDINGS CONTAINED IN THIS REPORT
ARE THOSE OF THE AUTHOR(S) AND SHOULD NOT BE CONSTRUED AS AN
OFFICIAL DEPARTMENT OF THE ARMY POSITION, POLICY, OR DECISION,
UNLESS SO DESIGNATED BY OTHER DOCUMENTATION.**

I. FOREWORD

The Fast Reduced Instruction Set Computer (F-RISC) project has been undertaken to explore the highest possible for computer clock rates using some of the most advanced semiconductor devices that have been developed in the US. The project originally capitalized on existing GaAs/AlGaAs Heterojunction Bipolar Transistors (HBTs) and microwave compatible Multi-chip Modules (MCM's) as the vehicles to achieve these goals. During this phase of the project the final architecture chips were fabricated at Rockwell, work began on the GE/HDI MCM package, a new thrust in SiGe HBT technology with IBM was started, and other advanced devices have been examined for even faster computer operation. The project can be expected to impact applications ranging from "super" workstations, and parallel processing nodes in TeraOPS/PetaOPS computers, to virtual reality engines for simulation, HDTV for high resolution imaging, media access controllers for fast microwave communication networks, and direct Digital Signal Processing (DSP) at high frequencies. These latter applications might be suitable for radar, high-speed encryption/decryption, and data compression/decompression. More importantly some of these device technologies might offer alternative directions should the evolution of CMOS as the primary microprocessor technology encounter difficulties.

The goal established for the (D)ARPA/ARO grants of the F-RISC series were originally to create a demonstration Fast RISC integer engine with a 2 GHz clock rate and a peak throughput of 1,000 MIPS. The project took 8 years to approach its final state. In the meantime CMOS computer technology advanced from 20 MHz clocks to 1500 MHz and even higher for experimental systems. It might seem at first blush that with the limited resources of a university design group this work would be rendered irrelevant by the much larger resources of industry. However, much has been learned from the GaAs 2 GHz clock effort. Moreover, HBT technology has also advanced. With faster and higher yielding SiGe HBT technology 16 GHz clock RISC engines now seem possible with modest optical lithography. Even 32 GHz seems remotely possible at room temperature if device/materials improvements continue with SiGe/C, and 64 GHz at Liquid Nitrogen Temperature (LNT). The purpose of a university program should be to explore new technology during its earliest phases in order to assess its impact, and to train students to cope with the problems of even faster designs. Hence even as the current program approaches its completion, this mission is still an important one. CMOS still faces significant hurdles in the future, and promising alternative/complementary device research tracks should be explored.

At the inception of the present work Rockwell International offered the Rensselaer team the opportunity to employ their 50 GHz baseline HBT process for this project. Typical gate delays for that HBT process were revealed by Rockwell to be approximately 25 picoseconds, and in spite of yield limitations, with reasonable pipelining it has been possible to create an architecture that could respond in about 10 gate delays per clock phase, or 250 picoseconds. Given the low initial yield expected with this process a multichip architecture rather than a monolithic single chip microprocessor was proposed. Typical chip yields of 20% at 5,000 HBTs were assumed for the purpose of the

demonstration originally, but this needed to be upgraded to 8,000 HBTs during the course of the project. Most of the additional devices were needed to make the chips testable at microwave frequencies using boundary scan based, embedded at-speed test circuitry.

Yields for the Rockwell 4 inch wafer line were limited primarily by Ga rich surface oval defects and problems with the gold-polyimide interconnection process. The Yields for the devices alone were limited by having an oval defect land in the emitter base junction or from emitter to collector where it will appear to be a short. The typical oval defect is about 1-10 microns in size and is much more conductive compared to the semi-insulating ideal GaAs material. The final run at Rockwell opted to use 2 micron by two micron square emitters. Just computing the probability that a defect of this size would actually land in the emitter area [6] one is led to the conclusion that device yields of about 60% are possible at 10,000 ovals per cm², but that this falls to 30% for 4000 HBT's at 500 ovals per cm². Other yields for the two level or three level metal gold lift-off process draw this down considerably. Anecdotally, such failures can usually be seen visually with an optical microscope. Often balls of gold or lift off detritus can be seen. Lift-off is of course needed since the gold is a noble metal and not readily etched. But the process of lifting off the excess gold after desired gold is inlaid into lift-off trenches is extremely error prone. But the gold metalization is still preferred to limit power supply droop, to enjoy electro-migration free operation, to improve skin effect losses, and for contact metallurgy with the GaAs system.

Fortunately, Rockwell's yields improved during the period of this project to meet this requirement. Yields as high as 50% for 10,000 HBT circuits are now possible at Rockwell. However, these require use of large 2-micron by 2-micron emitter openings. This has kept the emitter current high in FRISC/G, requiring the MCM to dissipate several hundred watts. The power dissipated by HBT circuits is typically set by the size of the area of the emitter that is possible with the device. Since the original emitter stripe specified by Rockwell was 1.4 microns by 3 microns, the minimum feature size of the device actually increased over the 8 years of the project from 1.4 microns to 2.0. The significance of this large emitter size and the implied large emitter current were perhaps not recognized at the onset of this project, and the resulting large power dissipations limit the ability to explore such a large device in more aggressive architectures. It is important to note that for bipolar design emitter size influences primarily this power dissipation, while the ultimate speed of circuits is still dominated by base transit time. So an important advantage of the HBT in CML is its ability to decouple power considerations from speed considerations. The smaller the emitter size, the lower the power. The thinner the base, the faster the device operation attained. SiGe HBT technology is now at 0.18 micron emitter size, and correspondingly the power levels have dropped by a factor of 10 relative to GaAs at 2-microns by 2-microns. Yields have also increased in SiGe to the point where several hundred thousand HBT's can be found in one chip at 20-30% yields. Unlike GaAs, SiGe exhibits no materials system specific defect such as the oval defect so that these devices in mass fabrication should evolve to at least 700,000-1,000,000 HBT's at reasonable yield numbers. The 700,000 number is compatible with what Exponential observed for a Si only bipolar process to obtain their 704 chip. These numbers stand to improve further as the devices are moved for production from the IBM East Fishkill

development line to the Burlington production line. IBM reports that there are no discernable SiGe defect changes that are related to the percentage Ge content of the base. If this is the case, then yields greater than 1 million SiGe HBT's should become possible in mass production. Significant CPU's can be constructed with this many devices, and if their clock rates are high enough then commercial applications may develop around this technology.

This will be seen as an advantage if the next generation of Fast RISC is pursued.

Specifically, as an option for this contract DARPA requested that we explore alternative SiGe HBT technology available at IBM for comparison. This HBT in SiGe appears to have evolved quickly from a Yorktown research project in the early 90's to a production technology today. The processing line identified as 5HP has an HBT with an f_T of 45 GHz and a f_{max} of 60 GHz. The minimum recommended emitter size for the SiGe HBT is only 0.5 microns by 1.0. This has led to approximately a reduction by a factor of 2.5 in emitter current from 2 mA to 0.8 mA. In addition the SiGe V_{be} turn-on voltage is 0.7 V or half that of the GaAs/AlGaAs device at 1.4 V. The result is a factor of about 5 reduction in power dissipation levels. The next generation of SiGe HBT would actually be twice as fast and achieve this speed in half again the power. Processing advances seem to have led quickly to a very aggressively downsized emitter in SiGe. The notable reduction in power and increase in yield provides a route to a single chip implementation of the Fast RISC, rather than the multi-chip solution used for GaAs.

An additional attractive feature of the SiGe HBT process at IBM is that it provides the possibility of co-integration with a contemporary version of CMOS with comparable lithographic feature sizes. One of the immediate advantages of this co-integration is the access provided to more advanced interconnection systems commonly used in CMOS. Furthermore, CMOS could provide a way to lower power dissipations further by replacing the hot HBT cores of memories in the present design with lower quiescent power circuitry. To continue to maintain speed the CMOS memory would have to access into extremely wide fields selectable by faster HBT decoders, something reminiscent of F-RISC/G L1-L2 memory 1024-bit transfer path. In addition, a one line lower level cache called L0 would provide the ultimate speed needed by the CPU. Most of the heat dissipated in the F-RISC/G is in the 16 L1 cache chips. Hence, overall reduction in power dissipation with a SiGe HBT BiCMOS design could be much larger than a factor of 5 in the present SiGe line (termed 5HP at IBM).

IBM also has discussed plans for its follow on "100 GHz" process in which the HBT emitter stripe may be as small as 0.18 microns by 0.5 microns. Based on this anticipated area reduction, the faster HBT would dissipate at least another factor of 2-4 less power than the 50 GHz HBT. IBM terms this line its 7HP process. In this process we can predict that 16 GHz clock operation would be possible at a power dissipation less than one 20th of the present system assuming only HBT devices were used. But since the majority of the FRISC/G power dissipation (75%) is in its L1 memory where CMOS

could be used with wide access and HBT decoding, a power dissipation even less than this is possible

As a result of the option, some pieces of the FRISC/G computer design were recast in the SiGe 50 GHz SiGe HBT baseline process at East Fishkill using one of the DARPA sponsored multi-user shared reticle runs under a contract to IBM using MOSIS as the broker. A 0.18 ps register file was implemented and verified in that process which was a full 32-bit by 32-word memory with triple porting for two read and one write port. This file, which was much larger than the file used in the lower yielding GaAs versions, is amenable to single chip RISC implementation. Furthermore with micro-pipelining used in the port address decoding, file accesses at 8 GHz appear possible. In this strategy a flip-flop is placed at the end of the file address decoder on all the word lines. The ongoing memory reads and writes then can take place while the next address is being decoded from the instruction decoder register. Since about half the file's delay is in its decoder, this permits the approximate doubling of the throughput of one of the key components of computer architecture. This style of micro-pipelined memory operation is due to Chappell and Chappell at IBM for fast cache access, but we have adopted it for use in the register file. The yield on these register files was excellent, essentially resulting in a working file on the first site touched down by the probe set. It is anticipated that when implemented in the new IBM 100+ GHz HBT process that 16 GHz operation of this file will be possible in this micro-pipelined mode.

For comparison, the most recent paper on a 0.25 micron CMOS register file at Yorktown Heights gave a read access time of 640 ps, and this was for a file with only 16 words of depth, though it was for a wider 64-bit word configuration. Allowing another gate delay for increased depth and doubling this time to scale backward to a comparable 0.5 micron lithography the hypothetical advantage of the bipolar design to the CMOS design would be 7.5 times faster operation. When the 100 GHz SiGe 0.25 micron HBT is available this would be the exact ratio. Of course, the relative ease of scaling of CMOS processing still puts the more difficult to implement HBT at a disadvantage. However, this situation may change as the minimum feature size approaches approximately 0.1 micron and optical lithography can no longer be used for CMOS.

Encouraged by the early success of the SiGe HBT register file another student, Matt Ernest, has developed a way to implement a 32-bit adder in only 5 CML macro delays. Assuming these delays to approach 13 ps, an adder could approach 67.5 ps or just shy of 16 GHz, as implemented in the 50 GHz IBM SiGe line. Taken then to the 100 GHz next generation it is clear that even 64-bit addition at 16 GHz would be possible using technology only a few years away from implementation at the present writing.

The early simulations for this register file and adder resulted in a proposal to DARPA in 1997 under BAA9703 for a 16 GHz advanced version of the F-RISC/G computer, called F-RISC/H which would be an enhanced F-RISC/G recast in this faster, higher yielding process. The machine would be designed in the form of a VLIW superscalar, with 8 processing fields, leading to a 16 GHz clock 128 GOPS (1/8 TeraOPS) machine suitable

for MCM implementation. The proposal received an award letter assigning the AO number F377 for the effort, but funds for this program have not yet materialized. This proposal remains our best recommendation for the future of room temperature supercomputing at this time. In the meantime work has been expanded to include design of fast internet chips using the SiGe HBT. Using IBM's 50 GHz SiGe HBT line Internet communication in the range of from 10 to 20 Gb/s is possible, depending on the method of clocking the data. This work complements the processor design effort well and provides yet another window with which to evaluate the HBT as a digital device for the future.

At the time of this report IBM has announced [EE Times, March 5, 2001] that the successor to the so called 5HP 45 GHz SiGe HBT 0.5 micron BiCMOS process will be a 120 GHz fT SiGe HBT BiCMOS process with both minimum emitter stripe and CMOS FET channel length (as drawn) of 0.18 microns. The SiGe HBT devices in this process are 2.67 times as fast as the earlier 5HP devices and the peak of the fT vs. Ic curve is at 200 uA for the minimum device, so the devices are both faster and cooler than those used in the GaAs project. The process still does not use Carbon in the base alloy.

Work summarized in this report is documented in its entirety on the DARPA mandated web site for the project,

<http://inp.cie.rpi.edu/research/mcdonald/frisc>

Details omitted from this current final summary report can be found in many of the theses created during sponsorship of the project on line at this web-site, which will be maintained for at least ten years subsequent to the termination of project funding. Thereafter the theses can be obtained in the traditional fashion from the Rensselaer library.

Table of Contents

I.	FOREWORD	iii
I.1	LIST OF FIGURES.....	x
II.	FINAL REPORT	1
II.1.	F-RISC / G OVERVIEW AND STATEMENT OF THE PROBLEM STUDIED	1
II.1.1.	HISTORICAL BACKGROUND AND MOTIVATION FOR THE RESEARCH	1
II.1.2.	COMPARISON WITH THE FUTURE OF CMOS	8
II.1.3.	OTHER MORE FUTURISTIC DEVICES	24
II.1.4.	SELECTION OF THE HBT FOR THE F-RISC.....	40
II.2.	SUMMARY OF RESEARCH RESULTS.....	42
II.2.1.	THE FAST RISC IN GAAS/ALGAAS HBT TECHNOLOGY.	42
II.2.2.	TECHNOLOGY	53
II.2.2.1.	CURRENT MODE LOGIC	53
II.2.2.2.	F-RISC / G CACHE IMPLEMENTATION	56
II.2.3.	ADVANCED MCM PACKAGING	57
II.2.4.	CACHE PIPELINE	64
II.2.4.1.	CACHE HIERARCHY.....	66
II.2.5.	MODEL DEVIATIONS AND THEIR IMPLICATIONS	76
II.2.5.1.	RE-IMPLEMENTATION WITH NEW MODELS	76
II.2.5.2.	QSIM LOGIC SIMULATOR	78
II.2.5.3.	TOP LEVEL FULL SPEED SIMULATION.....	81
II.2.5.4.	LOGIC PATH OPTIMIZATION	81
II.2.5.5.	DISTRIBUTION	84
II.2.5.6.	CLOCK DISTRIBUTION STRATEGIES.....	85
II.2.6.	FABRICATION OF THE CHIPS	89
II.2.7.	INSTRUCTION AND DATA CONFIGURATION.....	104
II.2.7.1.	CACHE CONTROLLER DESIGN.....	105
II.2.8.	COMMUNICATIONS.....	106
II.2.9.	INTRA-CACHE COMMUNICATIONS.....	108
II.2.10.	SECONDARY CACHE COMMUNICATIONS.....	110
II.2.11.	VIRTUAL MEMORY SUPPORT.....	111
II.2.12.	TIMING	112
II.2.12.1.	LOAD TIMING.....	113
II.2.13.	STORE TIMING.....	115
II.2.14.	INSTRUCTION FETCH TIMING	118
II.2.15.	OTHER CACHE STALLED	120
II.2.16.	PROCESSOR START-UP	120
III.	LIST OF ALL PUBLICATIONS AND TECHNICAL REPORTS	126

IV. LIST OF ALL SCIENTIFIC PERSONNEL SHOWING ADVANCED DEGREES EARNED BY THEM WHILE EMPLOYED ON THE PROJECT FROM ITS INCEPTION (1991).....	132
V. LIST OF INVENTIONS BY NAME	133
VI. REFERENCES.....	134
VII. BIBLIOGRAPHY.....	137
VIII. APPENDIX.....	141
REGISTER FILE / CACHE RAM OPTIMIZATION PROCESS.....	142
REGISTER FILE CIRCUIT SENSITIVITY ANALYSIS AND COMPONENT MODIFICATIONS.....	145
REGISTER FILE CIRCUIT MODIFICATIONS.....	150
REGISTER FILE OPTIMIZATION SUMMARY	156
IX. SAMPLE PUBLICATIONS	157

I.1 List of Figures

FIGURE 1. FORCES ON ELECTRONS AND HOLES IN THE VICINITY OF VARIATION OF THE BANDGAP.....	3
FIGURE 2. SIMPLIFIED CURRENT FLOW IN FORWARD BIASED HOMOJUNCTION BJT.	4
FIGURE 3. BAND DIAGRAM OF AN NPN TRANSISTOR WITH A WIDE GAP EMITTER SHOWING THE VARIOUS CURRENT COMPONENTS AND THE HOLE REPELLING EFFECT OF THE ADDITIONAL ENERGY GAP IN THE EMITTER (FROM KROEMER).....	6
FIGURE 4. CROSS SECTION OF ROCKWELL HBT (FROM ASBECK).	8
FIGURE 5. CROSS-SECTION OF FET (FROM MEINDL) SHOWING MESH INTO VOLUME ELEMENTS LARGE ENOUGH TO HOLD ONE DOPANT ATOM <i>ON THE AVERAGE</i> . STATISTICALLY, SOME OF THE VOLUME ELEMENTS WILL BE EMPTY IN ANY GIVEN DEVICE DUE TO PROBABILISTIC UNCERTAINTY.....	9
FIGURE 6. PLOT OF THE DISTRIBUTION DENSITY FUNCTION OF EFFECTIVE DOPING CONCENTRATIONS (FROM MEINDL) FOR AVERAGE DOPING DENSITIES OF $N_A = 5E17cm^{-3}$ AND $N_A = 3E18cm^{-3}$	15
FIGURE 7. MEINDL'S PLOT OF LEAKAGE CURRENT VS. CHANNEL LENGTH SHOWING NEARLY 7 ORDERS OF MAGNITUDE OF LEAKAGE INCREASE PER GATE AS THIS LENGTH DECREASES FROM 0.35 MICRONS TO 0.07 MICRONS FOR MEMORY, LOGIC LEAKAGE INCREASES ONLY SLIGHTLY OVER 4 ORDERS OF MAGNITUDE INCREASE, BUT WITH AN UNCERTAINTY OF NEARLY 4 ORDERS OF MAGNITUDE.....	17
FIGURE 8. DUAL GATE SOI CMOS DEVICE FOR N-CHANNEL FET SHOWING UPPER AND LOWER GATES.	18
FIGURE 9. THRESHOLDS FOR P+-P+, N+-N+ AND P+-N+ DOPED SILICIDE DUAL GATE SOI DEVICES (SUJII).....	19
FIGURE 10. SCHOTTKY BARRIER CMOS (SB-CMOS) EMPLOYING Pt OR Er SILICIDE SOURCE AND DRAIN MATERIALS TO FORM THESE BARRIERS. GATE INDUCED HIGH FIELDS AT THE TOP OF THE SOURCE NEAR THE GATE INDUCES TUNNELING THROUGH THE BARRIER (FROM TUCKER).	21
FIGURE 11. DEVICE TURN-OFF CHARACTERISTICS FOR THE SB-CMOS N-CHANNEL GATE AT $L=0.05$ MICRONS AND TOX OF 35\AA . CONVENTIONAL CMOS TYPICALLY EXHIBITS TWO OR MORE ORDERS OF MAGNITUDE GREATER TURN-OFF RATIO.	22
FIGURE 12. SAMPLE HETEROSTRUCTURE SiGe P-MOSFET. THE HEAVY DOPING LAYER IS LAYING HORIZONTALLY ALONG SIDE OF THE SiGe HOLE CHANNEL (FROM MEYERSON).	23
FIGURE 13. Pt GATE, STRAINED LAYER HIGH ELECTRON MOBILITY TRANSISTOR (HEMT) FROM (MEYERSON).	24
FIGURE 14. CROSS SECTIONAL VIEW OF THE DUAL ELECTRON LAYER TUNNELING TRANSISTOR (FROM SIMMONS).....	26
FIGURE 15. RETROGRADE DOPING PROFILE SHOWING INADVERTENT N_s DOPING AND INTENTIONAL N_A DOPING (FROM NING AND TAUR). THRESHOLD SHOULD BE SET BY HIGH N_A DOPING AND N_s SHOULD BE ZERO BUT MAY BE UNINTENTIONALLY DOPED.	27
FIGURE 16. VERTICAL FET STRUCTURE PROPOSED BY STURM.	29
FIGURE 17. TYPICAL SiGe HBT (STURM INCLUDED SMALL AMOUNTS OF CARBON IN THE SiGe ALLOY).....	30
FIGURE 18. INCREASE IN BREAKDOWN VOLTAGE WITH LIGHT CARBON ALLOYING IN BASE (FROM STURM). LEFT TRANSISTOR (A) USES $Si_{0.73}Ge_{0.25}$, RIGHT TRANSISTOR (B) USES $Si_{0.743}Ge_{0.25}C_{0.007}$	30
FIGURE 19. COMPARISON OF f_T VS. I_C PLOTS FOR VARIOUS EMITTER AREA PARAMETERS.	32
FIGURE 20. CURRENT FLOW BETWEEN EMITTER AND COLLECTOR THROUGH THE BASE, SHOWING LATERAL CURRENT PATHS THAT ARE GENERALLY MUCH LONGER THAN THOSE IN THE INTRINSIC BASE REGION. AS THE CURRENT DENSITY INCREASES IN THE BASE, THE LATERAL VAN DER ZIEL BASE PUSH-OUT EFFECT BECOMES MORE DOMINANT (FROM ROULSTON).	33
FIGURE 21. HITACHI SiGe HBT BASE SIDE CONTACT WITH REDUCED FRINGING CURRENT PATHS.	34
FIGURE 22. ZOOM ON HITACHI SiGe HBT SIDE CONTACT AT THE CORNER OF THE BASE REGION SHOWING THE NITRIDE GUARD RING AROUND THE INTRINSIC BASE. THE GAP FOR THE BASE CONTACT IS SHOWN AS 50NM OR 500 \AA , WHICH GREATLY RESTRICTS FRINGE CURRENT PATHS.	34

FIGURE 23. FT VS. IC CURVE FOR THE HITACHI BASE SIDE CONTACT TRANSISTOR SHOWN IN FIGURE 22 AND FIGURE 21. NOTE THAT THE PEAK IS AT ABOUT 150 MICRO-AMPS BUT HAS A LONG EMITTER STRIPE OF 1.7 MICRONS. THIS COULD BE REDUCED SOMEWHAT, PERHAPS TO 0.5 MICRONS WITH PROPORTIONATE REDUCTION IN CURRENT.	35
FIGURE 24. THRESHOLD SHARPENING AT LNT FOR SiGe HBT (FROM CRESSLER).	37
FIGURE 25. 33% FT PEAK IMPROVEMENTS AT LNT (FROM CRESSLER).	37
FIGURE 26. 100% IMPROVEMENT AT LIQUID NITROGEN TEMPERATURE ON FT VS. IC CURVES FOR 100 GHZ DEVICE (ZEROUNIAN).	38
FIGURE 27. YIELD CURVES MEASURED BY HITACHI ON A 10,000 TRANSISTOR RING OSCILLATOR TEST STRUCTURE OFF THEIR DEVELOPMENT LINE. ONE HALF MILLION BIPOLARS ARE POSSIBLE AT 20 % IN THIS DEVELOPMENTAL LINE. IN PRODUCTION YIELD SHOULD BE MUCH BETTER.	39
FIGURE 28. GAAS HBT YIELDS VS. THE MAJOR DEFECT (GA RICH OVALS) DENSITY (FROM ROCKWELL). COMPARE WITH THE HITACHI SiGe HBT YIELD IN FIGURE 27.	40
FIGURE 29. COLOR MICROGRAPH OF ROCKWELL GAAS/ALGAAS HBT. THE MINIMUM FEATURE SIZE IN THIS TRANSISTOR IS 1.4 MICRONS WHICH IS THE WIDTH OF THE CENTRAL EMITTER FINGER SHOWN CONNETED FROM THE TOP. BASE CONTACTS ARE ON EITHER SIDE OF THE EMITTER. THE COLLECTOR CONTACTS ARE FURTHER TO EITHER SIDE AND EXIT FROM THE BOTTOM OF THE FIGURE ON METAL 2. BASE CONTACT IS ON METAL 1. EMITTER CONTACT IS METAL 1.	44
FIGURE 30. COMPARISON OF THE ROCKWELL GAAS/ALGAAS HBT LAYOUT VS. MORE RECENT GENERATIONS OF SiGe HBT LAYOUTS.	45
FIGURE 31. THE FAST RISC ARCHITECTURE SHOWING SOME OF THE MORE CRITICAL AND LESS CRITICAL SIGNAL PATHS. THE SEVEN STAGE PIPELINE IS EVIDENT FROM THE PROGRAM COUNTER OR PC SHIFT REGISTER ON THE RIGHT. A FEED-FORWARD PATH IS IMPLEMENTED AROUND THE REGISTER FILE TO PREVENT HAVING TO WAIT FOR A WRITEBACK TO OCCUR WHEN THE NEXT INSTRUCTION REQUIRED THE SAME DATA.	46
FIGURE 32. COMPARISON OF RC WIRE CHARGING DELAYS ON M1, M2, AND M3 OF THE ROCKWELL 50 GHZ HBT BASELINE PROCESS. THE THREE T LINES CORRESPOND TO TRANSMISSION LINE BEHAVIOR FOR THREE DIELECTRIC CONSTANTS.	50
FIGURE 33. FAST RISC CACHING STRATEGY THROUGH L2 SHOWING 1024-BIT PAIR PATH BETWEEN L1 AND L2 FOR TRANSFERRING A SINGLE LINE OF CACHE BETWEEN L1 AND L2 IN ONE CYCLE.	52
FIGURE 34. DIFFERENTIAL CURRENT SWITCH.	53
FIGURE 35. SAMPLE STD. HEIGHT CML CELL FROM FRISC CELL LIBRARY.	54
FIGURE 36. F-RISC / G SYSTEM.	56
FIGURE 37. POSSIBLE MCM CHIP ARRANGEMENT WITH L1 AND L2 CHIPS IN PLACE.	58
FIGURE 38. CRITICAL PATH DIAGRAM.	58
FIGURE 39. DATA CACHE CRITICAL PATH.	59
FIGURE 40. ADDRESS TRANSFER FROM CPU TO CACHES.	61
FIGURE 41. SINGLE BUS ADDRESS TRANSFER FROM CONTROLLER TO RAMs.	61
FIGURE 42. DUAL BUS ADDRESS TRANSFER FROM CONTROLLER TO RAMs.	62
FIGURE 43. INSTRUCTION TRANSFER - RAM TO ID.	62
FIGURE 44. INSTRUCTION TRANSFER - RAM TO ID.	63
FIGURE 45. SEQUENTIAL CACHE OPERATION.	66
FIGURE 46. PIPELINED CACHE OPERATION.	66
FIGURE 47. SIMULATIONS: 2 KB HARVARD CACHES, DIRECT-MAPPED, SPICE TRACE.	67
FIGURE 48. SIMULATIONS: 2 KB HARVARD CACHES, DIRECT-MAPPED, TEX TRACE.	68
FIGURE 49. SIMULATIONS: 2 KB HARVARD CACHES, DIRECT MAPPED, GCC TRACE.	68
FIGURE 50. SIMULATION RESULTS FOR BENCHMARK SUITE.	69
FIGURE 51. 2 KB HARVARD CACHES, DIRECT-MAPPED, BLOCK SIZE EQUALS BUS WIDTH.	70
FIGURE 52. CPI AS A FUNCTION OF SET SIZE, BLOCK SIZE.	71
FIGURE 53. EFFECT OF SET SIZE.	72
FIGURE 54. EFFECT OF ARCHITECTURE ON CPI.	73
FIGURE 55. EFFECT OF ARCHITECTURE ON CPI.	74
FIGURE 56. EFFECT OF HARVARD CACHE SIZE ON CPI.	74
FIGURE 57. METAL GEOMETRIES FOR CAPACITANCE COMPENSATION STRATEGY, DUE TO UNEXPECTED ANISOTROPY IN ROCKWELL POLYIMIDE (WHICH INCREASES LATERAL CAPACITANCE).	78

FIGURE 58. REDUCING SIGNAL CROSS-TALK.	80
FIGURE 59. SOLVING CRITICAL PATHS WITH PIPELINE ADJUSTMENTS.....	82
FIGURE 60. SPLITTING A MASTER-SLAVE LATCH TO ALLEVIATE TIME CONSTRAINTS. A MASTER-SLAVE LATCH MAY BE SPLIT TO PROVIDE A SLOW LOGIC PATH WITH MORE TIME TO THE INPUT OF THE SLAVE LATCH.....	83
FIGURE 61. H-TREE CLOCK DISTRIBUTION. LOCAL CLOCK BUFFERS ARE DRIVEN BY A SYSTEM CLOCK BUFFER THROUGH A BALANCED H-TREE OF INTERCONNECT. CLOCK DISTRIBUTION IS GIVEN PRECEDENCE OVER ROUTING RESOURCES TO INSURE A BALANCED CLOCK NETWORK.....	86
FIGURE 62. LOCAL CLOCK DISTRIBUTION. AGGRESSIVE CLOCK DISTRIBUTION MAY INCREASE PERFORMANCE BY SKEWING THE CLOCK TO MATCH LOGIC PROPAGATION DELAY, HOWEVER THIS MAKES THE SYSTEM SENSITIVE TO HOLD TIME VIOLATIONS.....	87
FIGURE 63. CLOCK TRANSITION WITH RESPECT TO DATA. TO MINIMIZE POSSIBLE SETUP OR HOLD TIME VIOLATIONS, THE CLOCK MUST ARRIVE IN THE CENTER OF THE "SAFE TRANSITION" WINDOW.	88
FIGURE 64. FINAL F-RISC/G FABRICATED RETICLE, SHOWING THE FOUR MAIN ARCHITECTURE CHIP SITES. RIGHT STRIP OF SMALLER CHIPS CONTAINS ANOTHER COPY OF THE RPI TEST CHIP AND DESKEW TEST CHIP.	90
FIGURE 65. THE DATA PATH (DP) CHIP.....	91
FIGURE 66. COMPONENT OVERLAY FOR DATA PATH (DP) CHIP.	93
FIGURE 67. ZOOM ON FIGURE 66 SHOWING CLOSE-UP PHOTOMICROGRAPH OF 32W X 8B REGISTER FILE, REGISTER FILE DECODER ON LEFT AND SOME OF THE BOUNDARY SCAN SUPPORT CIRCUITRY, TOWARDS THE LEFT OF THE FILE.....	94
FIGURE 68. INSTRUCTION DECODER (ID) CHIP.....	95
FIGURE 69. SCHEMATIC OF INSTRUCTION DECODER CHIP.	96
FIGURE 70. L1 CACHE CONTROLLER (CC) CHIP. IDENTICAL CHIP USED WITH PERSONALIZATION FOR DATA AND INSTRUCTION CACHE.	97
FIGURE 71. ZOOM FOR FIGURE 70 IN THE VICINITY OF THE TAG RAM FILES.	98
FIGURE 72. SIMPLIFIED CACHE CONTROLLER BLOCK DIAGRAM.	100
FIGURE 73. FLOOR PLAN FOR CACHE CONTROLLER CHIP.	101
FIGURE 74. UNIVERSAL L1 DATA AND INSTRUCTION CACHE CHIP.....	102
FIGURE 75. CACHE RAM CHIP LAYOUT.	103
FIGURE 76. SCHEMATIC DIAGRAM OF CACHE MEMORY.	104
FIGURE 77. LOAD CRITICAL PATH COMPONENTS.....	106
FIGURE 78. COMPONENTS OF ADDER CRITICAL PATH (ADAPTED FROM [PHIL93]).	106
FIGURE 79. ABUS PARTITIONING.	108
FIGURE 80. DATA CACHE TIMING - CLEAN LOADS.	114
FIGURE 81. DATA CACHE TIMING - LOAD COPYBACK.....	116
FIGURE 82. INSTRUCTION CACHE MISS TIMING.....	117
FIGURE 83. DATA CACHE DURING INSTRUCTION CACHE STALL.....	119
FIGURE 84. INSTRUCTION CACHE DURING A DATA CACHE STALL.	120
FIGURE 85. INSTRUCTION CACHE AT START-UP.	122
FIGURE 86. INSTRUCTION CACHE DURING TRAP.	123
FIGURE 87. LOAD COPYBACK IN F-RISC / G CACHE.....	125
FIGURE 88. REGISTER FILE PERFORMANCE WITH VARIOUS DEVICE AND INTERCONNECT MODELS.	141
FIGURE 89. ACCESS TIME SENSITIVITY TO ADDRESS, BIT AND WORD LINE CAPACITANCE.	143
FIGURE 90. MEMORY CELL LAYOUTS.	145
FIGURE 91. WORDLINE SENSITIVITY TO TOTAL DECODER RESISTANCE.	146
FIGURE 92. WORDLINE AND ADDRESS-LINE CURRENT SENSITIVITY TO DECODER RESISTOR RATIO (TOTAL RESISTANCE = 440 OHMS).	147
FIGURE 93. ADDRESS LINE SENSITIVITY TO DECODER RESISTANCE RATIO.	148
FIGURE 94. READ/WRITE LOGIC BITLINE SWINGS DURING WRITE (PULL-UP RESISTANCE = 600 Ω).	149
FIGURE 95. ACCESS TIMES AND BITLINE SWING DURING WRITE FOR VARIOUS READ/WRITE LOGIC PULL-UP RESISTOR VALUES.	149
FIGURE 96. ACCESS TIME SENSITIVITY TO BITLINE CURRENT.....	150
FIGURE 97. INTERNAL SIGNAL SWINGS DUE TO RELATIVELY STATIC AND DYNAMIC ADDRESS CHANGES.	151

FIGURE 98. ORIGINAL AND REDUCED-WORDLINE-SWING MEMORY CELL DESIGNS.	152
FIGURE 99. "BRIDGE" RESISTOR BETWEEN BITLINES.	154
FIGURE 100. SENSITIVITIES TO BITLINE BRIDGE RESISTOR.	155
FIGURE 101. PERFORMANCE SENSITIVITY TO BITLINE BRIDGE RESISTOR.	155

I. FINAL REPORT

II.1. F-RISC / G Overview and Statement of the Problem Studied

II.1.1. Historical Background and Motivation for the Research

The mission of this project has been to explore alternative devices for fast computer design. In this phase of our contract work we have focused on the Heterojunction Bipolar Transistor or HBT as the candidate for this effort. This section of the report gives an overview of various candidates that were considered, and our reasons for pursuing the HBT.

One of the original inspirations for the work of this contract series was a Special Issue of the Proceedings of the I.E.E.E. published in January of 1982 on Very Fast Semiconductor Technology, edited by Richard Eden [1]. A key invited paper in that special issue was "Hetero-structure Bipolar Transistors and Integrated Circuits," by Herbert Kroemer [2]. The final sentence of the abstract for that paper was "the present overwhelming dominance of the compound semiconductor device field by FET's is likely to come to an end, with bipolar devices assuming at least an equal role, and very likely a leading one." Kroemer might well be expected to favor the device since he was an early proponent of it [3]. The idea of the HBT goes back to an early patent by Shockley in June 1948 [4].

By implication one might well have removed the term "compound semiconductors" in Kroemer's discussion of FET's since the MESFET in the compound semiconductor materials system (at least at low field strength) is faster than comparably scaled CMOS FET's. But of course, the evolution of the semiconductor industry increasingly into CMOS as a preferred medium for expression of microprocessor and memory design could perhaps only have been guessed in 1982. The good news was that CMOS only had to lithographically shrink to attain faster speeds, but the bad news is that it had to continue to shrink lithographically forever in order to keep delivering performance improvements.

The first big challenge will come around 0.1 microns minimum channel length, when conventional optical lithography runs into problems. At this point the industry will be forced to strike off into dramatically different lithography technology, and run the risks associated with that detour. This includes EUV, X-Ray, and Ebeam Projection technology. All of these depend on 1:1 projection schemes and require precise alignment over large 2 cm or greater reticle areas and very thin masks.

Anticipating this, in the summer of 1997 an article written by Bijan Divari, an IBM Vice President, the longevity of CMOS in this role was called into question. Essentially the year 2004 was called a "brick wall" to continued shrinkage of lithography and device size as the main evolutionary path for CMOS, and that "alternative device architectures" would have to be explored. The year 2004 is significant in that this is approximately the time frame for phase-out of optical lithography. IBM has been a vigorous explorer for

alternatives to optical lithography including x-ray, Extended UV or EUV, and Ebeam lithography. Whether the year 2004 is the correct date or not the key item is the reference to "alternative devices." Of course, alternative device architectures could include just SOI or surround-gate MOS, but hidden in the discussions that surround this interesting announcement are other worrisome signals that something much more aggressive might prove necessary. Even if the FET could continue to shrink using some new lithography the device itself might leak too much due to poor threshold control and excessive sub-threshold leakage or poor noise margin.

In recent past history, "Moore's Law" has described the trend line for advancement of computer technology of a performance doubling every 3 years or so. This law is not a law of physics, but rather a law of economic imperative. This is approximately the rate of advancement needed to command the venture capital needed to sustain the industry based on the size of the market it currently serves worldwide. Any disruption of this trend line due to physical barriers threatens the economic health of the industry, and with it any strategic vision based upon microelectronics and information technology. Unfortunately for a long time now it has been known that the trend line could not continue indefinitely. But the semiconductor fabrication process is so complex that technical predictions are notoriously nonspecific with regards to end-times for this technology. SEMATECH has posted a schedule called the National Technology Road Map, which assumes that devices of dimensions 0.025 microns could be fabricated and that progress could continue until the years 2012-2014. The plan did not identify or anticipate any disruption that could prevent the attainment of that goal. Scenarios deviating significantly from the model of CMOS shrinkage have not been considered in the devising of the SEMATECH model. Of course the consequences of such a derailment are extremely serious.

Later in the fall of 1999 Mark Packan of INTEL sounded a similar warning that continued advancement of CMOS might encounter many problems and that the semiconductor industry faced some of the most severe challenges in its history. The economic and political implications of a maturation of one of the highest of the High Tech industries was hardly noticed in the New York Times article summarizing the pronouncement. The scaling of HBT devices is significantly different from CMOS, and so one of the questions is whether this alternative device in particular has anything to offer in surmounting these challenges.

II.1.1.1. The HBT and Band Gap Engineering




Figure 1. Forces on electrons and holes in the vicinity of variation of the bandgap.

The earliest HBT's were fabricated in the III-V materials system because of the higher electron mobility found in that system. Processing difficulties for the HBT in that system has caused its evolution to lag behind the relative simplicity of fabrication in silicon for ordinary bipolar and CMOS. The financial resources of existing CMOS based fabrication also tend to favor continued advances in the same direction until some unexpected challenge or barrier proves insurmountable or uneconomical. Recent advances in SiGe technology, however, confer some of these advantages of Si-based technology back toward the HBT. This would appear to make it feasible to create very fast computing systems with optical lithography. This, of course, was always the promise of the bipolar device.

The Kroemer article focused on bandgap engineering because "the forces acting on the electrons and holes in a semiconductor are equal (except for sign in the case of electrons) are proportional to the slopes of the edge of the band in which the carriers reside". In several figures in the article (Figure 1 of which is reproduced here) Kroemer pointed out that this slope can be manipulated by varying alloy compositions in the various regions of the HBT. With wide bandgap emitters, electron injection from the emitter towards the collector could be independently manipulated from the hole injection from the base towards the emitter. The result would be a new way to increase the relative injection rates of these two carriers which results directly in higher beta, and in fact a beta which could be made high almost independent of the doping levels in the device.




Figure 2. Simplified current flow in forward biased homojunction BJT.

To first order, for the normal npn homo-junction BJT as shown in Figure 2, during forward bias conditions holes are swept towards the emitter from the base, and electrons travel from the emitter through the base to the collector where they are captured and the maximum value that the beta or current gain for the homo-junction transistor can attain is

Equation II.1-1

$$\beta_{\max} = \frac{I_c}{I_b} = \frac{D_n N_{de} X_e}{D_p N_{ab} W_b}$$

where the D's are diffusion constants for holes and electrons, the N's are doping densities and X and W are the distances through which the respective carrier travels. We can see that if the emitter doping is really large relative to the base, that the current gain can be large even if the ratios of the other quantities in this equation are unfavorable (as they often are). Intuitively this is because there would simply be many more electrons created by the high emitter doping than the number of holes generated by the base doping. However dopant levels are limited by solid solubility of the dopant in the host species, and so there is a limit to doping higher in the emitter, and often the base doping is sacrificed to achieve a high beta. But this then increases base resistance.

When the base emitter junction is forward biased the electrons from the larger number of donors in the emitter travel towards the collector, and the smaller number of holes travel towards the emitter, thus creating the current gain for which the bipolar transistor is well

known. However, this strategy of decreasing the base doping relative to the emitter for increasing beta comes at a price, because the amount of base spreading resistance increases. This is because inevitably there is an upper limit on emitter doping in a common homo-junction BJT, which results in a requirement to *lower* base doping to get the beta value up. But this adversely affects the ability of the device to charge and discharge parasitic capacitance in the device and from wire loading as will be discussed in the next paragraph.

Probably the best measure of switching time applicable to HBT digital circuits is the estimate of Dumke, Woodall, and Rideout (DWR) [5] who estimate the switching time as

Equation II.1-2

$$\tau_s = \frac{5}{2} R_b C_c + \frac{R_b}{R_L} \tau_b + (3C_c + C_L) R_L$$

where R_b and R_L are the base spreading and load resistance, C_c is the collector capacitance, C_L is the load capacitance, and τ_b is the base transit time. One can see that in two of the terms for the DWR formula the base resistance occurs. Hence low base resistance is important unless the last term is dominant. DMR examine the case of changing the load resistance to optimize this delay further finding that the optimum value of that resistance is

Equation I.1-3

$$R_L^{opt} = [R_b \tau_b / (3C_c + C_L)]^{1/2}$$

This optimum load value is often unrealistically low, but again shows the importance of the base spreading resistance in choosing better load values. Nevertheless, using this as the load resistance value the absolute minimum value of the switching time is

Equation I.1-4

$$\tau_s^{opt} = \frac{5}{2} R_b C_c + 2[(3C_c + C_L) R_b \tau_b]^{1/2}$$

Since the base spreading resistance figures so importantly in these formulas, strategies to minimize this parameter are key to capturing the ultimate speed of the device.




Figure 3. Band diagram of an npn transistor with a wide gap emitter showing the various current components and the hole repelling effect of the additional energy gap in the emitter (from Kroemer).

Figure 3 shows the band diagram of an npn transistor with a wide gap emitter showing the various current components and the hole repelling effect of the additional energy gap in the emitter. The emitter for the GaAs/AlGaAs HBT has a wider bandgap than the base as shown here.

The advantage of the wide bandgap emitter HBT vs. the simpler BJT is that the beta for the transistor is given by the following formula

Equation I.1-5

$$\beta_{\max} = \frac{I_c}{I_b} = \frac{D_n N_{de} X_e}{D_p N_{ab} W_b} \exp(\Delta E_g / kT)$$

where ΔE_g is the change in the energy gap between the emitter and base. Denote the mean speeds v_{nb} , and v_{pe} due to combined effects of drift and diffusion of electrons at the emitter end of the base and of holes at the base end of the emitter respectively, these two current densities become

Equation I.1-6

$$J_n = N_{De} v_{nb} \exp(-qV_n / kT)$$

and

Equation I.1-7

$$J_p = N_{Ab} v_{pe} \exp(-qV_p / kT)$$

where

Equation I.1-8

$$\Delta E_g = q(V_p - V_n)$$

Taking the ratios of the two current densities results in the modified formula for beta for the heterojunction. For a large positive change in the energy gap, ΔE_g , beta is significantly improved relative to the homojunction. However, in digital applications beta need not be over 100 to be effective. Hence, the improvements due to the exponential term are traded for a higher base doping, thereby reducing the base spreading resistance. Furthermore, even the emitter doping can be reduced somewhat so that emitter capacitance is usually ignorable.

The devices used in the F-RISC/G CPU are based on AlGaAs/GaAs heterojunction bipolar transistor technology as supplied by Rockwell International. ¹Figure 5 shows a cross-section of the Rockwell HBT device. The baseline process produces transistors with a nominal f_T of 50 GHz.

¹ As the details of Rockwell's HBT process and design rules are restricted by a nondisclosure agreement between Rockwell and Rensselaer Polytechnic Institute, only previously published information can be presented here.




Figure 4. Cross section of Rockwell HBT (from Asbeck).

As mentioned above, the primary advantage of using HBTs is that the heterojunction provides good emitter injection efficiency, lower base spreading resistance than in bipolar junction transistors (BJTs), and lower emitter-base capacitance (C_{je}). The GaAs / AlGaAs system also offers other advantages. Among them, electron mobility is high (on the order

of 5000-8000 $\frac{cm^2}{V \cdot s}$ in pure material vs. 800-2000 $\frac{cm^2}{V \cdot s}$ for Si.), reducing base transit time and charge storage at the emitter junction, and a semi-insulating substrate is available (on the order of $5 \times 10^8 \Omega \cdot cm$).

II.1.2. Comparison with the Future of CMOS

All of these advantages come with the higher base doping possible in the HBT. One of the most important implications of base doping was noted by Dr. Eden in his own overview article. In this article Eden noted that the heterojunction bipolar device had the best threshold uniformity of all the devices considered in the issue. This was obviously due to the fact that the turn on voltage for the device depended on the bandgap and not on doping. However he did point out that collector-emitter punch-through was dependent on doping. In his words, "with modest base doping levels in homo-junction bipolar transistors the number of doping atoms...becomes so small ($\sim 10^2$) in the base that simply Poisson statistical variations in the number of dopant atoms can lead to punch-through in a statistically significant number of transistors." The high doping level in the base for the HBT approaches two orders of magnitude higher levels (10^{19}) than found in homo-junction BJT's because it is still possible to have high current gain with the HBT. Hence, punch-through is significantly less likely to occur on random transistors in a circuit with many devices. Furthermore the threshold voltage for V_{be} at which significant bipolar collector emitter conduction occurs for the GaAs HBT is largely set by just bandgap and temperature and is as Kroemer states "very nearly a universal constant." The measured variation in the Rockwell GaAs HBT used in this project is only 3 mV out on an average of 1.3 V, barely 0.1% variation on a given wafer or between wafers [6].

Of course the comparable observation can be made for Si FET's. But here there is no "fix" for the problem. A state of the art CMOS FET has a channel length approximately 0.1 microns, and the channel width in full conduction is around one fifth of this dimension. For a minimal sized transistor with W/L ratio of unity, the total volume of the channel approaches that of a state of the art HBT, or about $0.2 \times 10^{-15} \text{ cm}^3$. Doping this region at 10^{17} would lead to only 50 dopant atoms and again statistical variations resulting from the Poisson process model for ion implantation would dictate a significant number of devices with even lower numbers of atoms and might be difficult to put into cutoff.

Meindl [7] in a very recent article has further quantified this and has shown that indeed this becomes one of the primary challenges faced by conventional CMOS because the threshold voltage V_{th} for turning on the device in the case of the FET depends on this doping level. Not only is the low number of dopant atoms a problem, but the placement of these dopant atoms within the channel becomes statistically non-uniform due to random variations in density. Meindl referred to this as an "intrinsic" problem because it depends only on mathematical probabilities. Keyes was apparently one of the earlier workers to recognize this would eventually become a problem. This lack of uniformity can lead to drastic increases in sub-threshold leakage current. For example, with only 50 dopant atoms on the average in the channel the standard deviation for a Poisson process is proportional to the square root of the average, and this is 7 atoms when the average is 50. Clearly when considering the additional requirement of spatial uniformity these requirements are even more stringent. For example if the channel is merely divided into 5 compartments with an average of 10 dopant atoms in each cell, the standard deviation is 3 dopant atoms, and there is significantly greater risk of not having any in a given cell. In a recent paper James D. Meindl and his students analyzed this effect, including examination of the uniformity issue by dividing up the channel into small cells.




Figure 5. Cross-section of FET (from Meindl) showing mesh into volume elements large enough to hold one dopant atom *on the average*. Statistically, some of the volume elements will be empty in any given device due to probabilistic uncertainty.

Meindl analyzed the FET by modeling it with a series of small MOS capacitor structures. He took each capacitor to be l^2 in area and X in depth, where X is the depletion depth at threshold and where the volume

Equation I.1-9

$$l^3 = \frac{1}{N_A}$$

on the average has exactly one impurity atom in it. Any specific elemental MOS capacitor will have instead m dopant atoms in it leading to the local dopant concentration in that volume of

Equation I.1-10

$$n_A = \frac{m}{l^2 X}$$

Under conditions outlined in the paper Meindl showed that the threshold voltage for each of the elemental MOS capacitors could be written as

Equation I.1-11

$$V_{th} = V_{fb} + 2\psi_B + \frac{1}{C_{ox}} \sqrt{\frac{2q\epsilon_s(\psi_B + \phi_f(m))m}{X}}$$

where again m is the number of dopant atoms.

Equation I.1-12

$$\phi_f(m) = \frac{kT}{q} \ln \left(\frac{m}{2l^2 X n_i} + \frac{1}{2} \sqrt{\left(\frac{m}{l^2 X n_i} \right)^2 + 4} \right)$$

and

Equation I.1-13

$$\psi_B = \frac{kT}{q} \ln \left(\frac{N_A + \sqrt{N_A^2 + 4n_i^2}}{2n_i} \right)$$

As pointed out by Meindl for a given gate bias, V_g , there exists value m_{\max} , such that the elemental MOS capacitor with that number of impurities or less will be inverted. The value of m_{\max} can be computed through these equations. Since m is a random variable the probability of finding an inverted elemental channel is

Equation I.1-14

$$p = \Pr(\text{inverted} | V_g = V_{th}) = \int_0^{m_{\max}} f(m, M) dm$$

where M is the average number of impurities in the volume, $f(m, M)$ is the probability density function for the number of dopant atoms in the elementary MOS capacitor, which is obviously a function of M . For the overall MOS FET to conduct electricity, there must be enough of the elemental MOS capacitors on to form a "percolation" path through the channel that is a region of such elemental volumes that form a continuous path. To do this the problem is broken up into two parts, the first to determine if there are I "inverted" MOS elemental capacitors out of K compartments. The probability for this is written $Y_K(I)$ and obviously is a Binomial probability distribution:

Equation I.1-15

$$Y_K(I) = \frac{K! [p^I (1-p)^{K-I}]}{I! (K-I)!}$$

Meindl's coworkers developed computer codes to determine whether if I elemental MOS capacitors function, they just happen to line up in such a way as to form a continuous channel path from source to drain. These simulations can be used to determine $Z_K(I)$, the probability there is a conductive path from source to drain. Since there are many ways for this to happen for many values of K, the total probability for finding continuous paths is the summation over all I

Equation I.1-16

$$Q(K) = \sum_{I=0}^K Y_K(I) Z_K(I)$$

Meindl approximated the distribution of ion implanted dopant densities with a Gaussian probability function, although for the low number of dopant atoms a more accurate distribution would have been a Poisson model. Use of the Gaussian model enabled him to obtain an explicit expression for p in terms of the average doping density N_A :

Equation I.1-17

$$p = \frac{\operatorname{erf}\left(\frac{n_A X N_A^{-2/3} - X N_A^{1/3}}{\sqrt{2 X N_A^{1/3}}}\right) + \operatorname{erf}\left(\sqrt{\frac{X N_A^{1/3}}{2}}\right)}{1 + \operatorname{erf}\left(\sqrt{\frac{X N_A^{1/3}}{2}}\right)}$$

where

Equation I.1-18

$$X = 2 \sqrt{\frac{\epsilon_s \psi_B}{q N_A}}$$

While p could be obtained in closed form, and hence in principle Y , the Z parameter had to be computed. Apparently the distribution of the ratio of I to K for channel continuity to occur appeared Gaussian such that if I/K were greater than about 0.59 chances were good that connection from source to drain would occur. The following formula was used to compute the "connection" probability given K cells with I in inversion for a square channel (whose width, W , equals its length, L):

Equation I.1-19

$$Z_K(I) = \int_{-\infty}^{((I/K - 0.59)/\alpha)} \exp(-t^2/2) dt$$

where

Equation I.1-20

$$\alpha = 0.149 \exp(-0.0636 L N_A^{1/3})$$

and L is the channel length. This would be the effective channel length, not the as-drawn channel length. Clearly as $L/l \rightarrow \infty$ connection approaches certainty, since the variance α goes to zero, and there are a large number of viable complete paths. Meindl used these expressions to compute what he termed the effective channel doping concentration $f(n_A)$ distribution density function among a large number of MOSFETs. This is the doping density required for a uniformly doped MOSFET to have the same threshold voltage. He used the chain rule to compute

Equation I.1-21

$$f(n_A) = \frac{dQ}{dn_A} = \frac{dQ}{dp} \cdot \frac{dp}{dm_{\max}} \cdot \frac{dm_{\max}}{dn_A}$$

leading to the following result for this probability density:

Equation I.1-22

$$f(n_A) = \sqrt{\frac{2}{\pi^2(\alpha^2 + \gamma^2)}} \cdot \left\{ \frac{(\epsilon_s \Psi_B)^{1/4}}{N_A^{13/12}} \right\} \cdot \left[\exp \left\{ -\frac{(p - 0.59)^2}{2(\alpha^2 + \gamma^2)} \right\} \right] \times \left[1 + \frac{\gamma(p - 0.59)}{(\alpha^2 + \gamma^2)} \left\{ \frac{1 - 2p}{2L^3 \sqrt{N_A} \sqrt{p(p-1)}} \right\} \right] \times \left[\frac{\exp \left\{ -\frac{(n_A X N_A^{-2/3} - X N_A^{1/3})^2}{2 X N_A^{1/3}} \right\}}{1 + \operatorname{erf} \left(\sqrt{\frac{X N_A^{1/3}}{2}} \right)} \right]$$

where

$$\gamma = \sqrt{p(1-p)} / L N_A^{1/3}$$

Using these expressions Meindl has an explicit closed-form expression for the probability density for the effective channel doping concentration for a FET with effective values for $W=L=0.07$ microns. Using this formula these coworkers were able to compute and plot this distribution for two target values of $N_A = 5E17 cm^{-3}$ and $N_A = 3E18 cm^{-3}$. We include this plot here to press the point. Obviously the higher the number of dopant atoms on the average, the larger the spread, but the variance depends on this spread divided by the average value and is worse for the smaller number. Roughly there are 50 atoms in the lower doped channel.




Figure 6. Plot of the distribution density function of effective doping concentrations (from Meindl) for average doping densities of $N_A = 5E17cm^{-3}$ and $N_A = 3E18cm^{-3}$.

Although the spread at the higher doping density seems wider, compared to the average, the spread is actually narrower normalized to that average value. Hence, it is the lower doped transistor that exhibits the wider relative spread. For the lower doping density, which is preferred in high-speed devices, the average would be 50 dopant atoms. But the value of the probability density is only down by 6 orders of magnitude at 250 atoms, indicating that million transistor circuits might already have many transistors at significantly different thresholds, and 100 million transistor circuits would surely have serious yield problems due to degraded noise margins, and leakage would climb. Meindl continues the analysis to show the actual probability distribution for the threshold voltages themselves, but one can see even from this analysis just how marginal the setting of the threshold can be.

Meindl analyzed only these mathematical uncertainties and termed them “intrinsic” limitations because they are mathematically present even if the fabrication of the device is perfect. In addition to these intrinsic limitations there are statistical fluctuations in threshold due to fabrication tolerance uncertainties. Meindl predicted that this intrinsic uncertainty in threshold values would become a yield limitation on fabricating CMOS at 0.07 microns effective channel length since the threshold uncertainty is so large that some

transistors would have threshold deviations of several hundred mV. In addition, sub-threshold leakage current was seen to escalate 6 orders of magnitude moving from 0.35 microns and 0.07 microns. This would prevent the evolution to smaller supply voltages required to keep the power dissipation density reasonable. Devices with thresholds on the low side would be more difficult to be driven into cutoff, leading to excessive leakage current with excessive DC power dissipation and possible incorrect output voltages since the logic threshold shifts with the device threshold. Devices with thresholds on the high side would operate slower than the design target. Once again logic thresholds could be shifted with the device threshold resulting in static errors. Of all the effects of threshold uncertainty perhaps none is more pervasive than the excessive power dissipation problem because while only a few actual logic errors might be expected in the form of another yield detractor, leakage will be the more prevalent effect. Subthreshold leakage can already be anticipated to increase with decreasing channel length, but shifting thresholds can dramatically increase this dissipation. We can see the degree of jeopardy this places traditional microprocessors in when we examine the DEC/Compaq Alpha Road Map. The most aggressive projected device generation at 0.125 microns (Leff of 0.1 microns) shows an anticipated power dissipation of 150 watts.

SHIP DATE	JUNE 1996	DEC. 1999	2Q 1999	2Q 2000		
	EV4	EV5	EV5	EV6	EV7	EV8
LITHO. (μM)	0.35	0.35	0.28	0.18	0.18	0.125
VDD IN VOLTS	2.5	2.2	2	1.5	1.5	1.2
PACKAGING	WP/PGA	WP/PGA	WP/PGA	FC/SCP	FC/SCP	FC/MCP
PINOUT	499	587	587	587	1443	1800
CLOCK IN MHZ	400	575	750	1000	1250	1650
SPECINT95	16.8	32	45	60	75	155
SPECp95	28	61	67	100	160	400
L2 BW (GB/s)	0.5	2	2	2	10	10
L1 BW (GB/s)	1.3	4	4.5	5.5	16	64
POWER (W)	55	95	80	60	125	150
DIE SIZE (MM ²)	210	320	210	125	400	350

DEC/Compaq Alpha Road Map [from their web site]

It is quite clear that power levels are becoming extremely high from this trend even without the aforementioned threshold uncertainty induced leakage power dissipation problem. What is needed is a device technology where DC currents decrease with shrinking geometries, not increase.

Semiconductor workers (both in Japan and the United States) have observed this "atomistic" limitation on conventional CMOS. The HBT like the BJT, of course, has its turn on voltage set by bandgap. Hence there is nothing comparable to the threshold uncertainty seen with the conventional bulk CMOS devices that sets an "intrinsic" limitation on threshold uniformity. Both device thresholds on the other hand are a function of temperature but here we have at least a chance for uniformity by proper thermal engineering.




Figure 7. Meindl's plot of leakage current vs. channel length showing nearly 7 orders of magnitude of leakage increase per gate as this length decreases from 0.35 microns to 0.07 microns for memory, logic leakage increases only slightly over 4 orders of magnitude increase, but with an uncertainty of nearly 4 orders of magnitude.

There is a natural effort to circumvent the problems of low numbers of dopant atoms in CMOS by eliminating the doping as the technique for setting thresholds for the FET. This has resulted in proposals for using work functions at the gate oxide semiconductor interface to accomplish this. One approach is the so-called dual gate approach, which uses heavily doped silicide gate material on the top and bottom of the channel. TI has pioneered this approach and has demonstrated it at 0.18 microns. However the device

structure begins to look much more complex using the top and bottom dual gate device. There are a couple of ways to try to fabricate these devices. One which has been tried is the bonded silicon wafer approach. First one wafer is processed to include the bottom connection and gate and the thin oxide layer. This is polished using Chem Mechanical Planarization. Then a second wafer is placed on top of this and bonded to it using atomic force bonding. The wafers need to be very flat. Next the top wafer is polished to only the thickness of the channel, or roughly 100 Å. Then the top gate is fabricated.




Figure 8. Dual Gate SOI CMOS device for N-channel FET showing upper and lower gates.

The dual gate SOI is required to set thresholds when the intended power supply voltage is very low. The threshold set for n channel FET is -0.1 V for N+-N+ dual silicided gates and 1.0 V for P+-P+ silicided gates. However, Suzuki and Sugii [4, 5] using the N+-P+ combination more usable average between these two levels is obtained. The threshold voltage for the dual gate structure is dominantly set by the work function of the gate material and is less sensitive to the actual doping levels in the channel. Hence, this strategy has been developed to head off the problems Meindl analyzed. Fujitsu has published a number of papers on SOI threshold setting by varying gate materials and doping levels. In this discussion the gate material is heavily doped polysilicon with electrode connections being made out of refractory (high temperature) Cobalt Silicide. Also, some conductors such as TiN have been discussed for doped gate material.

The work function is varied by heavily doping the gate material itself rather than the channel, and because the doping is so high, it is less sensitive to low dopant numbers in the channel. The polarity of doping in the upper and lower gate material can be varied with some schemes using P+ for both the upper gate, or lower gate, or N+ for both. The threshold for P+ is 1 V, whereas for N+ the threshold is -0.1 V, neither of which is adequate for high-speed, or for low-power. The only alternative is to mix these two strategies. In that situation the upper and lower polysilicon gate materials are differently

doped as P+ and N+. This structure results in two partial turn-on thresholds, the first of which occurs on the N doped gate interface to the channel and the second at the P doped interface. The net effect of these two turn-on voltages combines to set a kind of overall threshold. Note that the real action is that there are two effective FETs in parallel, one on the back side and one on the front, which combine to give a net threshold. The resulting threshold plot vs. thickness of the SOI channel is shown in the following figure:




Figure 9. Thresholds for P+-P+, N+-N+ and P+-N+ Doped Silicide Dual Gate SOI Devices (Sujii).

Here we can see that if the SOI thickness is small enough the effective threshold for the P+-N+ heavily doped dual gate configuration can place a threshold at about 0.25 V. This requires a SOI thickness of 100 Å, which may not be very consistent with good transconductance for the device, which requires high gate capacitance, and even thinner oxides at small device feature sizes. Furthermore the threshold is varying with this thickness, which requires excellent thickness control, or once again the threshold uncertainty will become noticeable again.

In addition to these sensitivities the dual gate structure is awkward to fabricate. First a wafer must be prepared that contains all the fine lithography patterned gate and electrode materials up to and including the thin oxide for the thin lower oxide. This must be polished to exquisite levels of flatness in planarization. Then another intact wafer must be bonded to this preexisting structure. The second wafer is then polished using chemical-mechanical polishing (CMP) to precisely the desired thickness of the channel (100 Å based on Figure 10). It is this extreme uniformity, which must be maintained for threshold control. Then the top gate must be aligned to the bottom gate with great precision. Given that there is no self-alignment possible between the two gates, yield may be expected to suffer. Whether the yield will be worse than what would have been experienced with dopant uncertainty depends on how short the channel length is.

Another recently announced process at IBM involves regrowth of recrystallized polysilicon over oxide so the lower gate oxide. Some of the techniques can involve using the raw polycrystalline thin film to using techniques to recrystallize the polysilicon while portions of the melt contact a crystalline substrate area. In such cases the silicon regrown by this approach can lead to some extension of the crystal structure over a part of the gate oxide. It is not clear at this point whether the yield of the bonded silicon CMP approach or the recrystallized Si or thin film polysilicon channel is better than limited dopant count conventional FET. In other words, conventional CMOS may have its lifetime extended if circuits are made with transistors different from those used at present. But this would require a large reversal of the present computer design paradigm focus, which is oriented towards highly parallel architectures with extremely large transistor counts leveraging from the high yields possible in CMOS. But by becoming overly dependent on these large transistor counts, architectures become boxed into a strategy, which might not progress. In short it would be prudent for the present to concentrate on conservative architectures with only modest number of transistors to insure continuity during the period of this evolution.

In short, the dual gate SOI FET device possesses most of the complexity of the HBT with none of its advantages. The structure is no longer a totally planar technology, but epitaxial growth cannot be employed to improve junctions between dissimilar materials, and there is no self-alignment possible to insure feature abutment. Current flow is still horizontal in the ultra thin channel. This limits the amount of current that can flow, and limits its transconductance. In the HBT current flow is vertical and the minimal device can readily handle many mA of current. The higher transconductance of the HBT makes it a superior device for charging primarily capacitive loads. The ultimate FET speed limit is still dictated by expensive lithographically defined gate length, as opposed to epitaxial defined base thickness for the HBT. Layer thicknesses can be greater for all regions of the HBT than they are for the FET, which may require gate oxides much thinner than 100 Å.

Another somewhat different idea for coping with the low numbers of dopant atoms in short channels was developed at HP by Snyder, Helms and Nishi, and discussed independently by Tucker. In this case the inventors tried to build into the source and drain a Schottky barrier effect using metal silicides for the drain and source materials. Metal silicides form natural Schottky barriers on silicon. They have low resistivity, and lend themselves to good contact resistance at small dimensions. Earlier efforts published in 1968 by Lepselter, and Sze on the same subject involved much longer gates and the right effects were not observed. However, at about 0.1 microns the effect can be used. The gate in these devices induces a high electric field near the top of the source barrier, which generates field emission Fowler Nordheim tunneling through the Schottky barrier into the channel. As noted by Tucker, Platinum Silicide and Erbium Silicide have Schottky voltages such that nearly symmetrical CMOS thresholds can be arranged at about a third of a volt, with a supply voltage of ~1.5 V. The following figure is taken from Tucker's paper showing the use of Cobalt Silicide for inter-device connection material.




Figure 10. Schottky Barrier CMOS (SB-CMOS) employing Pt or Er Silicide source and drain materials to form these barriers. Gate induced high fields at the top of the source near the gate induces tunneling through the barrier (from Tucker).

Experimental confirmation of the viability of these structures is not yet available. However Tucker was able to use Monte Carlo simulation to obtain V-I characteristics for the gate drain and source drain connections. Although the source drain characteristics appear normal it is the behavior around the threshold that appears unsatisfactory. The device does not seem to shut off as aggressively at room temperature as a more conventional FET. That is the device looks more like a slightly varying resistor than a transistor switch. More significantly the amount of current that can flow when the transistor is on is reduced due to the small region of initial conduction at the top of the gate. This figure shows a current level of only 100 micro-amps for a gate with of 1 micron, or a width to length ratio of 20:1. Tucker points out this effect could be mitigated by decreasing the oxide thickness further, or employing other silicides, but obtaining large on-currents is an inherent problem for shrinkage of a device whose vertical dimension fixes the current possible per unit length of the device. Tucker points out that the current due to tunneling is confined horizontally to flow from the top 2 nm of the source electrode. Since the bipolar transistor carries this current in a vertical direction through the device's thinnest region (the base) it can deliver a much higher current for a given size device when called upon to do so. Ultimate performance of circuits using these devices will depend not just on their intrinsic speed, which requires continual shrinkage, but also on the ability of the device to swing voltages on large the large capacitances due to fanout loads and interconnection parasitics. This requires that the device be able to deliver high current when necessary for these situations.




Figure 11. Device turn-off characteristics for the SB-CMOS n-channel gate at $L=0.05$ microns and T_{ox} of 35 \AA . Conventional CMOS typically exhibits two or more orders of magnitude greater turn-off ratio.

Another device track for the Si based FET is to combine it in heterostructure with SiGe alloy layers. These multilayer structures can implement concepts normally associated with GaAs/AlGaAs MESFET technology [10]. These include approaches such as presenting carrier injection to undoped channels (which then exhibit high carrier mobility due to the absence of scattering) from parallel layers containing dopant atoms at relatively high concentration. Interlayer strain in these structures can also be engineered to enhance mobility, leading to high f_T devices with superior transconductance. One example is the SiGe p-MOSFET presented by Verdonckt-Vandebroek in which a SiGe hole channel is deposited above a heavily doped Boron Modulation doping layer with a spacer layer of undoped silicon. The doping of the donor layer can be at a much higher level since carrier scattering in the actual channel is low. This provides one way to get around the low dopant count in conventional CMOS. At the same time the channel has no doping in it so there is low carrier scattering and improved mobility.




Figure 12. Sample Heterostructure SiGe p-MOSFET. The heavy doping layer is laying horizontally along side of the SiGe hole channel (from Meyerson).

Another possible contender for CMOS replacement is the Pt or Pt Silicided gate HEMT or high electron mobility transistor. In this case [10] SiGe is used to introduce tensile stress in an epitaxially grown silicon layer near the SiGe layer. This induced stress greatly increases the electron mobility in the channel. The use of a Schottky barrier gate eliminates the need for extremely thin oxide. Peak transition time frequencies of 140 GHz have been measured with a 0.1 micron gate. Au-Sb source and drain contacts keep the contact resistance down.

This type of HEMT structure was developed during approximately at the same time as the HBT in the SiGe materials system. However, when it came time for IBM to deploy a new device into its production line it was not the HEMT, but rather the HBT. One must speculate as to why this was chosen at the time. Evidently the benefit/complexity trade off favored the HBT. IBM had many years of bipolar fabrication capability and experience. So it was perhaps natural to favor it. There are also obviously fewer epitaxial layers involved. Each layer requires a lot of process control to optimize. In some ways the bipolar device then has come to be viewed as simpler than many of its competing FET device alternatives.




Figure 13. Pt Gate, Strained Layer High Electron Mobility Transistor (HEMT)
from (Meyerson).

Each of the CMOS alternatives that have been discussed carries with it a set of problems. Unfortunately there is risk associated with assuming that CMOS will develop along one of these lines without disruption. In the meantime HBT technology seems to be addressing several of these issues.

II.1.3. Other More Futuristic Devices

This discussion has ignored certain III-V MESFET related devices, which potentially can lead to fast circuit design. However, in past reports design experiments carried out with GaAs/AlGaAs H-MESFET SBFL at 0.7 microns minimum feature size to implement the core of the F-RISC led to disappointing circuit performance (400 MHz clock) relative to the GaAs/AlGaAs HBT design that remains the focus of the present (2 GHz clock) effort. In fact the performance of CMOS has already overrun that performance level. The conclusion is that the MESFET is less effective in driving interconnections and suffers more from the effect of wiring in architectures than the HBT, and at the same time does not seem to progress in lithography fast enough to outpace conventional CMOS. Consequently we will not attempt to survey recent developments in III-V MESFET technology. However, there remain possible materials systems in which the leverage of mobility advantages could bring these back into the spotlight. These include the InP/InGaAs/AlGaAs system, GaP and GaN. Our discussion of alternate devices for use in fast RISC has dwelt primarily on bipolar and FET types of devices. In recent years new devices have begun to emerge which to one extent or another employ tunneling as the primary transport mechanism for switching. These include the Josephson Junction (JJ), Resonant Tunneling Diodes (RTD's) and Resonant Tunneling Transistors (RTT). Tunneling offers extremely fast switching, but it also requires fabrication of extremely thin layers through which the tunneling can take place. The RTD and RTT involve the

use of companion heterojunction bipolar transistor which is considered slower than the tunneling devices. Logic can be computed with the resonant quantum effect portion of the device. But the bipolar device then plays a role similar to the totem pole bipolar driver in conventional BiCMOS providing the ability to drive large wire loads.

The Josephson Junction device offers switching times approaching a TeraHz, but at the moment demands 4 degree Kelvin operating temperature to permit Nb interconnect to exist in a super-conducting state. The JJ offers extremely low power supplies and low voltage swings in the range of a few 10^{-3} of mV. While this is attractive, the cost and complexity of operating a liquid Helium cryogenic system is equally unattractive. Past commercial efforts in this area at liquid Nitrogen temperatures were singularly unsuccessful. In addition, to attain the largest switching speed the junction must be extremely thin ($\sim 10\text{-}20$ Å). Current technology involves growth and patterning of Aluminum Oxide as this dielectric, on top of the Nb thin film. This is considerably more complicated than growth of comparably thin oxide layers thermally from a silicon wafer. Since this is not an epitaxial growth process the JJ suffers from somewhat the same yield limitations as dual gate SOI by the oxide growth method.

A related class of fast devices combines the idea of fast tunneling with two dimensional electron gas concepts. These require two "sheets" of electron occupancy layers that face each other and are separated by small thicknesses (~ 100 Angstroms) of wide bandgap semi-conducting layers. These must be double heterojunction (DH) structures to provide the band bending required to create the two dimensional electron gas sheets. In one implementation developed at Sandia Laboratories by Jerry Simmons [11], the two electron sheets have controllable tunneling through the use of a third electrode resembling the gate on a classical FET. This type of device also exhibits very high switching speed and employs a very low power supply voltage, making it extremely thermally efficient. Since the gap between the two electron gas sheets is fabricated by epitaxy the thickness of the insulating gap can be controlled in manufacturing more accurately than the thickness of grown oxides or other deposited insulators. This suggests that high yield might be possible assuming that other GaAs yield limiters can be controlled. Examples of such defects are known from HBT technology and include oval defects, which could short out the device between the two conducting sheets. It is also possible that this type of structure could be realized in the SiGe alloy system.




Figure 14. Cross sectional view of the Dual Electron Layer Tunneling Transistor (from Simmons).

Note that similar to bipolar device operation the controlled current travels vertically through the tunneling layer providing short tunneling distances leading to fast switching. However in this device the current must approach the region of tunneling through the extremely thin 2D electron gas sheets. Current enters and exits the device through source-drain connections on the ends with depletion gates to isolate the appropriate connections to the two dimensional electron gas layers, and must distribute through the top and bottom layers through these thin conducting sheets. These by necessity would be more lossy than passing the current directly in the vertical region to the collector as in the case of the HBT, even though the mobility of carriers in the 2D electron sheets might be quite high.

The area of the control electrodes on top and bottom of the tunnel layer determine the controlled current. The device combines some features of a FET and bipolar device operation, but does not rely on doping to set thresholds. Although currently the device is useful primarily only at 77 degrees Kelvin, its inventors believe that room temperature operation will prove feasible with a suitably modified device. The lower the temperature, the more tolerant the device fabrication is because the gap need not be as small. Thresholds for tunneling devices are extremely sensitive to thickness of the gap.

Another approach to dealing with the V_{th} uncertainty problem is the idea of a retrograde doping layer below the channel while using no doping atoms in the actual channel. This "proximity" effect depends on the depth below the channel surface boundary location, X_s , and works well when the doping in the retrograde region is very high, supposedly reducing the dopant atom uncertainty in that region. Taur [12] gives the threshold uncertainty as

Equation II.1-23

$$\sigma_{vth} = \frac{q}{Cox} \sqrt{\frac{NaW_{dm}^0}{3LW}} \left(1 - \frac{Xs}{W_{dm}^0} \right)^{3/2}$$

The relative uncertainty is this number normalized to the actual threshold voltage, which of course would be quite high, were it not for the second term. By controlling the proximity distance, Xs , one can reduce this. In fact, one can argue that by setting Xs exactly to the depletion distance this uncertainty can be made zero. One problem with this is controlling the size of Xs in a manufacturing environment. Since the coefficient in front of the second term is so high, a small error in controlling Xs gets magnified quickly. The second problem with this approach is assuring that the dopant atom concentration in the actual channel would now be precisely zero. Assuming that this is not zero but some Ns , while the retrograde concentration is Na .




Figure 15. Retrograde Doping Profile showing inadvertent Ns doping and intentional Na doping (from Ning and Taur). Threshold should be set by high Na doping and Ns should be zero but may be unintentionally doped.

The threshold voltage becomes:

Equation II.1-24

$$V_t = V_{fb} + 2\psi_B + \frac{1}{C_{ox}} \sqrt{2\epsilon_{Si} q N_a \left(2\psi_B - \frac{q(N_s - N_a)X_s^2}{2\epsilon_{Si}} \right)} + \frac{q(N_s - N_a)X_s}{C_{ox}}$$

or

Equation II.1-25

$$V_t = V_{fb} + 2\psi_B + \frac{qN_a W_{dm}^0}{C_{ox}} + \frac{q(N_s - N_a)X_s}{C_{ox}}$$

This shows that if N_s is non-zero, any uncertainty in this number of dopant atoms is superimposed on top of that for N_a . Although the intent is for N_s to be zero, inevitably some out-diffusion will occur in processing, or if the retrograde region is done by implantation, due to random stopping profiles. In the limit where X_s is exactly set to the maximum depletion width the third and fifth terms in this expression will cancel exactly and we obtain:

Equation II.1-26

$$V_{th} = V_{fb} + 2\psi_B + \frac{q(N_s)W_{dm}^0}{C_{ox}}$$

Which is precisely what we would have obtained if there were no retrograde well at all and N_s were the channel doping level. Since the intent would be to make N_s zero, the result of having N_s in fact be small but non-zero, would leave us with a threshold uncertainty that could be no better than that of a conventional FET. Uncertainties in W_{dm}^0 or C_{ox} simply make matters worse.

Sturm has published another interesting FET structure. This one is a vertical FET in which channel length would be fabricated epitaxially. This does capture some of the advantages of an HBT replacing the base as the device control region with a vertical channel. Although the issue of how to grow both n and p channel FET's in this strategy is never addressed, this approach would permit channel definition without the dependency on lithography. In addition, a novel use of SiGe:C permits use of another effect in which the presence of a small amount of C in the SiGe alloy would suppress boron out-diffusion

from p-doped regions for the device. This could be used in the source and drain regions for a p-channel vertical FET as shown in this case. In addition to the excellent control over boron dopant position distribution in the device the channel is totally surrounded by the gate. In the surround gate configuration total depletion of the channel is possible, and excellent sub-threshold leakage is obtained. Additionally substrate injection current is absent in this type of device and this makes the device a bit less noisy during switching. However to suppress punch through the channel was doped at 10^{18} , so once again in a small volume threshold uncertainty would be a problem, and in this configuration there would be no way to create the equivalent of a retrograde region. The growth of the oxide on the sidewall of the surround gate vertical FET, and the need to obtain large area contact vias to the source and drain also represents process and yield challenges. As with the upper-lower dual gate SOI device the processing steps seem more complex than for a HBT primarily because of the oxide growth step in comparison to simple junction or contact growth.




Figure 16. Vertical FET Structure proposed by Sturm.

However, the idea of suppressing boron out-diffusion was also exploited by Sturm in another paper in what appears to be a more appropriate setting, namely for thin base

formation in the npn HBT where Boron is also the dopant. In this case not only was the out-diffusion




Figure 17. Typical SiGe HBT (Sturm included small amounts of Carbon in the SiGe Alloy).




Figure 18. Increase in Breakdown Voltage with light Carbon alloying in base (from Sturm). Left transistor (a) uses $Si_{0.73}Ge_{0.25}$, right transistor (b) uses $Si_{0.743}Ge_{0.25}C_{0.007}$.

Examination of Figure 18, taken from the paper by Sturm [14], shows that the introduction of a very tiny atomic fraction of Carbon not only blocks out diffusion of Boron from the base, making thinner bases, and hence higher f_T 's possible, but the breakdown voltage BV_{ceo} is clearly pushed up. The beta has gone down, but that is consistent with having an elevated base doping which has not been optimized. Evidently the same mechanism that slows or blocks diffusion paths for Boron, also helps suppress electron avalanche in the vicinity of the collector. Thinning the base improves f_T by the formula

Equation II.1-27

$$\frac{1}{2\pi f_T} = \left(\frac{kT}{qI_c} \right) (C_{jc} + C_{je}) + C_{dCB}(r_e + r_c) + t_{bb}$$

where, ignoring Base-Collector Depletion Layer transit time, and Base-Emitter Depletion Layer transit times, and emitter transit times, and for a pure Si base

Equation II.1-28

$$t_{bb} \approx \frac{W_B^2}{2D_{nB}}$$

From the quadratic dependence on base width it can be seen that relatively small improvements in base thinning make large improvements in reduced base transit time. Ning and Taur derive that for the SiGe transistor (ignoring for the moment the small Carbon alloy content, that this improvement would be further amplified by the standard factor for this type of transistor:

Equation II.1-29

$$\left(\frac{t_{bb}(SiGe)}{t_{bb}(Si)} \right) = \frac{2kT}{\Delta E_{g,SiGe}} \left\{ 1 - \frac{kT}{\Delta E_{g,SiGe}} \left[1 - \exp\left(\frac{-\Delta E_{g,SiGe}}{kT} \right) \right] \right\}$$

All other formulas for transistor current gain, Early Voltage and actual collector current are similarly amplified. For a total base bandgap narrowing of 100 meV a low current base transit time of a SiGe base transistor is about 2.5 times smaller than that for a pure Si base. To repeat, this formula ignores effects of other components of transit time such as emitter crossing time and depletion crossing times, but these are generally also small for the SiGe transistor.

It can be seen from Equation II.1-27 that the f_T vs. I_c curve generally rises for increasing I_c until the base transit time dominates the expression, at which point the f_T saturates.

However, various effects at high current such as the Kirk base forward push-out effect and thermal effects cause the curve to eventually fall at very high currents. In other words there is little incentive to push the transistor past the point where the transit time dominates. Hence the ultimate speed is seen to be limited by the base transit time as expected, and operation of circuits should be limited to the low I_c (rising) part of the curve.




Figure 19. Comparison of f_T vs. I_c plots for various Emitter AREA parameters.

The issue of power dissipation for various size emitters is captured in Figure 19 for the IBM 5HP process. It can be seen as the AREA of the emitter is decreased the peak of the current at which the peak of the f_T vs. I_c curve is located is decreasing. Since in CML this current flows through the conducting path through the current tree, the current required for realization of (lightly loaded) fast gate delays, declines and with it the power dissipated at DC by the tree (currents in CML trees are DC for the most part). This justifies the quest for smaller emitter sizes in the HBT. However, it can be seen that as the emitter length approaches the emitter width the height of the peak is falling.

One source of the difficulty is the fringing currents that flow laterally around the emitter through the base to the collector. This is sometimes referred to as the Van der Ziel lateral base push out effect. Some feeling for the problem can be visualized from this figure from Roulston's book [15]:




Figure 20. Current flow between emitter and collector through the base, showing lateral current paths that are generally much longer than those in the intrinsic base region. As the current density increases in the base, the lateral Van der Ziel base push-out effect becomes more dominant (From Roulston).

It can be seen that as the emitter becomes smaller the lateral current paths which are slower due to longer path lengths, begin to dominate the total current flow. That is as L becomes smaller the paths in the edge region, L' , becomes more important in determining the effective transit time. Clearly the edge or fringe current paths are less important in larger emitters. But we seek small emitters to limit current flows. The smaller the device is, the lower the device parasitics are, and the lower the current is required to overcome these parasitics.

Hitachi has produced an interesting innovation designed to prune out these lateral paths. This is shown in the following figure:




Figure 21. Hitachi SiGe HBT base side contact with reduced fringing current paths.

Hitachi has explored a SiGe HBT with a Silicon Nitride layer surrounding the intrinsic base region, considerably restricting the lateral base push-out current paths. The primary reason for the region where the push-out effect occurs is because of the conventional contact strategy for the base and the tendency to just grow the extrinsic base on collector material. The nitride layer intercepts this interface preventing the “bleeder” fringe paths that increase f_T . Access to the base is provided by lateral vertical contacts at the corners or edges of the intrinsic base. A close up view of one of these corners is shown below




Figure 22. Zoom on Hitachi SiGe HBT side contact at the corner of the base region showing the Nitride guard ring around the intrinsic base. The gap for the base contact is shown as 50 nm or 500 Å, which greatly restricts fringe current paths.

The implications of the base side contact and the use of the nitride barrier layer are very great, because the emitter AREA scaling can continue to be scaled downward. In one plot Hitachi shows the f_T vs. I_C curve peaking at 45 GHz with an I_C of only 150 micro-Amps. But the emitter stripe is 0.2 microns by 1.7 microns. The length can clearly be scaled further downwards. Some degree of aspect ratio is still required to keep base spreading resistance down. Nevertheless one can imagine this being scaled down by at least a factor of 3 to peak at 50 micro-Amps. More significantly the current required to switch at half speed would be reduced by at least 4 in this case. The reason is that the curvature approaching the peak provides a wide arch over which power may be reduced while performance is reduced to a lesser degree less since the slope is flat at the top of the curve.




Figure 23. f_T vs. I_C curve for the Hitachi base side contact transistor shown in Figure 22 and Figure 21. Note that the peak is at about 150 micro-Amps but has a long emitter stripe of 1.7 microns. This could be reduced somewhat, perhaps to 0.5 microns with proportionate reduction in current.

Predicting what happens to the slope of the fT vs. I_c curve with future scaling is dependent on the scaling of the capacitances in Equation II.1-27, which the current I_c must overcome before reaching the peak. With base width shrinking the intrinsic base resistance will decline tending to make I_c higher, but by Equation II.1-28 if the base resistance is linear in base, width the improvement in t_{bb} is quadratic. We have seen this type of linear/quadratic tradeoff in CMOS voltage scaling.

Ning and Taur based on work by Solomon and Tang [16] have derived a key scaling parameter set based on several assumptions. One is that the voltage swing in CML/ECL would remain the same, and the power supply would remain the same. Another is that due to thinning of the base, base doping would have to climb with the inverse base width squared to prevent punch through. A full base width halving is not needed for doubling the peak fT value as has been already mentioned, but every other generation this could result in a halving of the beta.

Wire lengths would shrink by the scale factor and hence in the fringe field limit so would the wire capacitance. With fixed voltage swing the wire switching time scaling linearly by the scale factor would require fixed current from generation to generation. The smaller device design rules that go with such a scaling would then demand large current densities rising with the square of the scale factor. But such scaling strategies ignore other trends such as increasing the number of wiring layers or dropping the dielectric constant in the interlayer dielectric at the same time and use of Chemical-Mechanical Planarization (CMP) which can make thicker dielectric layers reducing capacitance further than mere lateral scaling would permit. Other tricks like 3D chip stacking may be needed to remove wire from digital architecture to keep the current density down in the emitter. Additionally many sections of logic contain very short wires that contribute negligibly to circuit delays. In these sections, current can scale downward with each generation. Only a few inter block wires on a chip are very long, and these could be shortened or have wire capacitance dropped by ILD or 3D strategies. For the few lines that remain, current would not scale as suggested by Tang and Solomon, but again the number of these wires might not be very great. Current scaling in the dominant logic core could continue to go down with each generation. However, in certain wire dominated sections of logic, such as memories and register files this current scaling problem may emerge as one of the key impediments to practical scaling.

Reduction of voltage swing remains another intriguing possibility for reducing power and/or increasing speed. But this appears to require lower operating temperature. The extreme of lowered temperature is Liquid Nitrogen Temperature (LNT), and sustained operating temperature would require extremely low power dissipation. Alpha EV-8 processors nominally operating at 150 watts would not be a particularly attractive combination for LNT operation. Cressler [17-18] has studied operation of the SiGe HBT at LNT. We note that GaAs HBT and Si homo-junction bipolar transistors do not operate well at LNT due to carrier freeze-out. However, Cressler has found them satisfactory, and in fact to operate better. In particular the threshold characteristics of the device turn

on are sharpened, raising the possibility of lower voltage swing. In addition Cressler found that the peak of the f_T vs. I_c curve was improved.




Figure 24. Threshold Sharpening at LNT for SiGe HBT (from Cressler).




Figure 25. 33% f_T Peak Improvements at LNT (from Cressler).




Fig. 1 Transist and maximum oscillation frequencies against collector current between 300 and 77 K at $V_{CB} = 0.8$ V

▲ 77 K
 ◆ 150 K
 ▼ 220 K
 ● 296 K

Figure 26. 100% Improvement at liquid nitrogen temperature on f_T vs. I_c curves for 100 GHz device (Zerounian).

The last evidence of LNT improvements comes from a paper by Zerounian, which shows that a device with a f_T at room temperature of 120 GHz and an f_{max} of 70 GHz at room temperature, improves to 213 GHz for f_T at LNT and 100 GHz for f_{max} .

In summary, many of Kroemer's arguments about the HBT remain intact after nearly two decades, with perhaps a couple of remaining notable exceptions. One of these is due to the difficulty of shrinking the emitter stripe feature size. This is largely a fabrication issue, but it puts a lower bound on the amount of power dissipated by the HBT that can be unacceptably high in GaAs technology.




Figure 27. Yield Curves measured by Hitachi on a 10,000 transistor ring oscillator test structure off their development line. One half million bipolars are possible at 20 % in this developmental line. In production yield should be much better.

The other issue is yield, which is also intimately related to the minimum feature size that can be built. Kroemer undoubtedly was thinking of the HBT as a III-V materials system device because at the time of writing of his paper SiGe did not exist. However, the real surprise may be that the HBT's concepts may be most successful in the SiGe alloy materials system. The SiGe system is one in which early aggressive minimum feature size shrinkage has had requisite processing resources to develop successfully. In SiGe yield appears to be much larger and co-integration with CMOS is not only possible but an accomplished fact. In an unusual publication Hitachi published actual yield information on a developmental line that individual HBT yield was at 99.9997% as measured in 10,000 bipolar ring oscillator circuits. These numbers are anecdotally not even as high as IBM reported yields for 40,000 HBT ring oscillators in a development line. By just squaring the number 0.999997 one can compute the yield for each doubling of the number of HBT's. At 18 squaring operations one has 20% yield, and $2^{19} = 524,288$ HBT's. This calculation is very sensitive. To accurately measure this yield would require building a circuit a lot closer to this size and measuring the yield for that. More than likely the yield in a real production line would be in the neighborhood of several million. Modern microprocessors such as the PowerPC3 require yield on 15 million transistors but only about a third to a quarter of these transistors are used for logic. The rest are for memory. A 32-bit integer RISC architecture requires only about 25,000 HBT's.




Figure 28. GaAs HBT yields vs. the Major Defect (Ga Rich Ovals) Density (from Rockwell). Compare with the Hitachi SiGe HBT yield in Figure 27.

Clearly with the defect density of GaAs HBT's a multichip module solution would be required for even a 32-bit Fast RISC demonstrator project.

II.1.4. Selection of the HBT for the F-RISC

The F-RISC/G GaAs/AlGaAs HBT project was the first to explore computer architecture with the HBT in this materials system in the so-called emitters-up orientation. TI in the mid 1980's explored HBT technology in the collectors up configuration using IIL technology. In IIL technology the emitters can be tied together in a common conducting sheet which results in greater circuit density but the HBT in the emitters down configuration did not seem as fast as the emitters up candidates. In 1991 at the onset of the DARPA HPCC Fast RISC project at Rensselaer the only available candidate for the HBT was the Rockwell 50 GHz baseline process. Gate delays of 25 ps seemed the fastest available at the time. A shadow project, called F-RISC/I was supported by IBM to test a 0.7 micron Rockwell H-MESFET process for comparison.

The minimum feature size (the emitter stripe) of Rockwell HBT device is not optimal for creation of such machines because the current required by the large emitter stripes is too high. To obtain usable yield the initial emitter stripe of 3 microns by 1.4 microns did not shrink as time evolved as expected, but, in fact, increased to 2 microns by 2 microns by the project's completion due to yield considerations. Nevertheless, the GaAs/AlGaAs system is not going to be limited to such feature sizes forever, and as these features are shrunk, power levels will decline. The GaAs/AlGaAs system still has on record some of the highest f_T values, with published results exceeding 200 GHz. At the time of the writing of the Kroemer paper the tables presented no supplier for HBT IC fabrication. In the intervening years Rockwell, TRW and TI all created such foundries with financial help from DARPA. Now, finally we engage the first real test of the GaAs HBT as a vehicle to attain high clock rates for computing, thanks also to DARPA/ARO sponsorship.

Near term SiGe appears to offer the next best a vehicle for probing the impact of HBT devices on computing because its feature sizes have shrunk faster, and its turn on voltage is only half that of the GaAs/AlGaAs HBT. Its yield is higher, power is lower, and at this point in history seems poised to offer an alternative to pure CMOS computer evolution, while providing co integration with CMOS. This co integration will permit some power savings in memory circuits. In the closing sections of this report we will provide a brief look at what may be possible in the very near future with SiGe HBT's.

II.2. Summary of Research Results

II.2.1. The Fast RISC in GaAs/AlGaAs HBT Technology.

Another key inspiration for the F-RISC project occurred during a visit by Robert Sherburne who accepted a teaching position in 1985 at Rensselaer after receiving his PhD on the DARPA sponsored project "Berkeley RISC II." Reduced Instruction Set Computers or RISC's derive their high clock rates using simple, streamlined architectures, which result in compact layout, minimal wiring effects, and favorable speed power tradeoffs. Hardware software tradeoffs are also scrutinized carefully using probabilistic models of the use of various resources. In effect hardware is added only if in a real mix of software applications the benefit is clearly established. Otherwise operations are performed with software. In addition, interest grew for using the RISC architecture as a vehicle to explore new, fast devices in higher performance computers implemented in initially low yielding technologies. It was evident even at the time the Fast RISC (F-RISC) project was initiated, that devices many times faster than the CMOS employed in the Berkeley predecessor project existed. But as shown in Figure 29, GaAs HBT IC's in the emitters up configuration have very low yield. This is partially due to the vertical topography of the device, partially due to oval defects inherent in the materials system, and partially due to the Au-polyimide "lift-off" process used for interconnections. Consequently RISC architecture would be the best vehicle for exploration.

At the time DARPA had sponsored projects at TI to examine a GaAs/AlGaAs HBT process using emitters down configuration for IIL circuit implementation of a MIPS processor. Other work was sponsored at RCA and McDonnell-Douglas to explore MESFET processor implementation. None of these processors achieved greater than about 120 MHz in speed clock speed. In the case of IIL very likely the low speed was due to the absence of emitter follower wire and load driving characteristics. While the MESFET enjoyed higher electron mobility, the III-V fabrication technology lagged CMOS in lithographic feature sizes. CMOS quickly overran these efforts. These past efforts point out the risks of over stating the physical or financial problems faced by CMOS, or underestimating the ability of CMOS to continue evolving on the Moore's Law projection. However, as we have noted earlier in the report there are also risks associated with assuming that this projection will continue unabated forever. At some point one or more of the limitations of CMOS will act to terminate or significantly slow down its evolution creating economic chaos in an industry that depends on rapid growth for high rates of financial return to investors. From a world strategic point of view a defense scheme dependent on holding the high ground in such technology would then be disrupted.

Work in the planning stages involved research in yield enhancing packaging strategies which included the idea developed at Rensselaer and Honeywell of Wafer Scale Hybrid Packages (WSHP), later referred to as Multi-Chip Module packaging (MCMs). These

packages were designed to densify low yielding integrated circuits in a format that approximated Wafer Scale Integration. Because wiring in such assemblies approximated that found on a chip, low yielding integrated circuits could be packaged to look like a high yielding single IC. This early work was sponsored by GE/CRD, which eventually productized the idea in the GE/HDI process. Several papers were published on this concept. However, a suitable host project to explore the MCM had to be identified.

Finding advanced devices in a cooperative proved difficult for the F-RISC group at first. Ted Creedon, a Rensselaer alumnus who was working at Tektronix in Portland, OR helped provide support for the project's first PhD candidate, Hans Greub. Ted provided not only recommendations for Tektronix to fund fellowships for Hans, but also provided a key contact with workers inside the company familiar with an early polysilicon emitter bipolar process with an f_T of 12 GHz. The first effort to gauge the speed of F-RISC was conducted by Hans while developing a cell library for this Tektronix process. As a part of the thesis Hans was able to create a few of the major building blocks for a fast RISC and from it estimate the speed possible, implying a clock rate of 250 MHz. The Tektronix foundry was still in its infancy when Hans made his first fabrication run. It yielded some ring oscillators, which confirmed the speed of basic gates. But there was insufficient yield at that point to observe register files working at speed. Tektronix's mission in instrumentation did not have a goal of attaining the extremely high yields needed for microprocessor production.

Obtaining serious funding to explore higher clock rates took nearly a half-decade. In 1990 John Toole at DARPA identified this project as one with promise and provided the initial funding for exploring the HBT as the device of choice for high speed. In addition, Art Cappon and Ken Elliot then at Rockwell met the principle investigator at a DARPA sponsored meeting reviewing DARPA's support for their GaAs/AlGaAs HBT effort. These individuals helped the F-RISC group to select Rockwell as the cooperating partner. The enthusiasm of both Ken and Art were for the project, helped cement a strong working relationship between Rensselaer and Rockwell for the next decade. The success of the F-RISC project has been in part due to Rockwell's patience and support through difficult technical challenges.

The initial funding for the F-RISC project was obtained as a result of the DARPA call for proposals for Strategic Computing Initiatives. Additional funding was obtained under the High Performance Computing Initiative (HPCC). Hans Greub, just completing his thesis under Tektronix fellowship support agreed to join the group as a research assistant professor, and became an extremely strong contributor to the GaAs HBT effort. Hans' original design presented in the thesis (called FRISC/E) did not employ a 4-phase clock. Hence one of the major changes in moving towards the GaAs/AlGaAs effort was a shift to increased pipelining using a 4 phase 2 GHz clock.




Figure 29. Color Micrograph of Rockwell GaAs/AlGaAs HBT. The minimum feature size in this transistor is 1.4 microns which is the width of the central emitter finger shown conneted from the top. Base contacts are on either side of the emitter. The collector contacts are further to either side and exit from the bottom of the figure on Metal 2. Base contact is on Metal 1. Emitter contact is Metal 1.

Figure 29 shows an optical photomicrograph of the actual HBT used in the project. With a minimum feature size of 1.4 microns an ordinary optical microscope can be used to photograph this picture. For comparison, the present (5HP) and next (7HP) generations of SiGe HBT's are shown in Figure 30. It can be seen that the device size for HBT technology has followed its own shrinking design rule road map. Of course a minimal CMOS device would be even smaller than the smallest of these devices.

GaAs and SiGe HBT Minimum Emitter Device Size Comparisons




Figure 30. Comparison of the Rockwell GaAs/AlGaAs HBT layout vs. more recent generations of SiGe HBT layouts.

The architecture of the Fast RISC is shown in Figure 31 here below. For anyone familiar with the Berkeley RISC and Stanford MIPS projects the blend of the two can be seen at

results from the level of pipelining involved. The second important area of modeling was the area of L1 cache stalls and transfer of lines of cache between L1 and L2. The simplicity of the architecture also permitted designers to focus on a modest number of interconnection tuning issues to tune the architecture in layout to achieve the desired speed. The doctoral thesis of Bob Philhower contains most of the initial architecture of F-RISC/G except for one timing problem solved by Steve Carlough just before final tapeout.

To increase the computer's speed the F-RISC/G demonstration engine implementation contains a seven-stage pipeline as shown in Table II.2.1-1. The I1, I2, D1, D2, and DW stages are all dedicated to memory accesses.

I1	Instruction Fetch 1	Transfer instruction address to cache on branch
I2	Instruction Fetch 2	Receive instruction from cache
DE	Decode	Decode instruction
EX	Execute	Execute instruction
D1	Data Read 1	Transfer data address
D2	Data Read 2	Receive data from cache if a LOAD
DW	Data Write	Cache modified if STORE; write result into register

Table I.2.1-1. F-RISC / G Seven-stage pipeline (Adapted From [Philhower]).

Like many modern RISC implementations, F-RISC relies on deep pipelines and a cache hierarchy to achieve high throughput. Table I.2.1-2 enumerates a number of RISC processors contemporary with F-RISC/G along with F-RISC, and illustrates some of the key architectural features of these processors.

The use of cache memory hierarchies has become paramount in computer design. Irrespective of the expense of massive amounts of extremely high-speed memory, packaging technology has not yet evolved to the point where the entire main memory of a high-speed computer can be placed in close enough proximity to the core CPU to allow reasonable access times. In fact this trend is likely to continue to get worse, almost to the extent that even L1 cache could not keep up with the processor clock, and a form of primitive one or two line cache, called L0, would be required to match processor clock rates. This push demands a concept called multi-threading, which is not implemented in FRISC/G, but was proposed in 1997 for the FRISC/H effort.

	Year	Bits	Clock (MHz)	Pipe Stages	Primary Cache (i/d)	
					KB	Assoc.
UltraSPARC ²	1995	64	167	9	16/16	1/1
Alpha 21164 ³	1994	64	300+	7	8/8	1/1
PA-RISC 7200	1994	32	140	5	—	Na/64
PowerPC 620	1995	64	130	4	32/32	8/8
MIPS R10000	1995	64	200	5	32/32	2/2
F-RISC/G	1995	32	2000	7	2/2	1/1

Table II.2.1-2. Architectural features of contemporary RISC processors.

Table I.2.1-3 lists some of the design details of the processors shown in Table I.2.1-2.

	Technology	Power (W)	Size (mm ²)	Devices
UltraSPARC	0.5 μ m CMOS	28	315	5,200,000
Alpha 21164	0.5 / 0.4 μ m CMOS	50	299	9,300,000
PA-RISC 7200	0.55 μ m CMOS	30	210	1,260,000
PowerPC 620	0.5 μ m CMOS	30	311	6,900,000
MIPS R10000	0.5 μ m CMOS	30	298	5,900,000
F-RISC/G	1.2 μ m GaAs HBT	250	10000 (MCM)	200,000

Table II.2.1-3. Contemporary RISC Processor Technology.

As shown in the table, the F-RISC / G core CPU and primary cache alone are expected to dissipate 250 W (or 2.5 W / cm²), illustrating the obvious problems that would be associated with packing even more memory onto the multi-chip module (MCM). This power dissipation is primarily due to the result of the rather large emitter stripe AREA for the GaAs/AlGaAs HBT used for this phase of research, which makes the emitter current excessively high. However, the rather low yield of the chips makes the power density

tolerable for MCM implementation. This part of the technology can be expected to improve totally in a SiGe HBT implementation.

One can observe in Figure 30 that there are only two core building blocks that set the ultimate limit on speed. One is the register file, and the other is the adder. Most of the rest of the building blocks are merely holding registers or multiplexers. The speeds of these building blocks were verified in the predecessor contract on one of the HSCD shared foundry runs at Rockwell. The read access time of the register file was verified at 213 ps, and the add time at 750 ps or three clock pipe cycles as seen in Table II.2.1-2.

However, as was to become one of the "discoveries" of the project the main determinant to clock speed turns out to be the wiring. The individual building blocks themselves involve very little internal wiring (mostly device capacitive loading) but the wiring in between the major blocks can be quite lengthy. It should be evident that processor speed cannot exceed the charge and discharge speed of the longest wire in the architecture, unless that wire is not needed to settle within one instruction period. The effect of wiring on speed of a processor should not be underestimated. It became the entire focus of the thesis by Steven Carlough. In that thesis extensive wire capacitance and resistance modeling had to be undertaken, with repeater amplifiers inserted in some of the more resistive. In addition, at that generation of the design, a third level of metal, M3 was introduced, which permitted better power distribution to prevent on-chip droop, and provide some lower resistance wiring paths for signals. The earlier designs had used only two gold metalization layers, which was consistent with the early mission of GaAs in analog applications. However as Internet and wireless applications grew, even the analog applications demanded some digital circuits and hence needed more layers of wiring.

Examining Figure 33 we can see that the RC charging time for a line 1 centimeter long on Metal layer 3 is only about 50% longer than the one way transmission line propagation time for a wire with an interlayer dielectric constant of 4. However, it is about twice that for an interlayer dielectric constant of 2 and nearly three times that for a dielectric constant of 1. An air dielectric constant insulator wire on M3 halves the unloaded propagation delay of approximately 50 picoseconds to 23 in 1 centimeter of travel. M2 wiring is slightly worse than M3 due to excess delay involving wire resistance in addition to source resistance. Observe that the curves for M3 and M2 are approximately linear with increasing line length while M1 is approximately quadratic in its rate of increase. Asymptotically the wires approach the linear delay performance of a transmission line as the resistance in the wire and circuit go to zero. One could say that the development time of the F-RISC/G would have been cut in half if there had been four levels of metal. The two level metal scheme left at least one direction of routing travel in the most resistive level of metal since the wiring in these layers is orthogonal in alternating levels. The legacy of this wiring delay used on multi millimeter paths proved too onerous to completely eliminate in later years of rework. As a result of this experience Rensselaer was a large proponent and major participant of the SRC CAIST (Center for Interconnection Science and Technology) whose goal it was to promote lower dielectric constant polymeric insulators, thicker insulators by Chemical Mechanical

Polishing or CMP, more and thicker layers of metalization, and lower resistivity metals for chip wiring.




Figure 32. Comparison of RC wire charging delays on M1, M2, and M3 of the Rockwell 50 GHz HBT Baseline process. The three T lines correspond to transmission line behavior for three dielectric constants.

As a result a gate driving a one centimeter run of wire on M1 is nearly 10 times slower than an unloaded gate, and nearly 5 times slower than a transmission line with interlayer dielectric constant of unity, which represents the ultimate in wire propagation speed. This excess resistance problem becomes noticeable for shorter wired driven by bipolar transistor drivers when compared with CMOS drivers because the drive impedance is lower for the bipolar case. A bipolar driver with a 100 ohm drive impedance starts to exhibit quadratic delay loading behavior when the wire resistance reaches this same level of 100 ohms. This always happens on the lower levels of the interconnection stack first because wire dimensions are most aggressively shrunk there. In modern integrated circuits there is a "reverse" scaling trend as the level in the stack ascends. This is for two reasons. The first is improved yield, and the second is lowered resistance. The larger cross sectional area used on the upper levels of the stack makes it easier to fabricate these layers, and less prone to defects. The typical philosophy is to "protect" the "investment" in the lower fabrication steps by going to higher yielding steps toward the end of the

stack. This also insures that the upper levels on the top of the stack can provide "faster" wiring there. Hence these levels are used for the critical signal path or longer wire connections whose rise/fall time bandwidth needs to be optimized.

One of the implications of this philosophy is that modern microprocessors that obtain their higher speed through device shrinkage need to retain or increase their scale on upper levels of wiring. Since this decreases the capacity for any of these upper layers to hold wire, one result is that each successive generation of microprocessor needs more layers than the previous generation. Since each successive generation decreases the available area quadratically with the scale factor, whereas the wire length is decreasing linearly, the net result is a rapidly increasing number of layers per generation. Whereas three layers of wiring were needed for 2 micron design, 4 were needed by 0.8 micron, 6 by 0.35 micron, and now 8 for 0.22 microns. As one approaches 0.13 micron geometries this number of layers would exceed 10. With two lithographic and processing steps per layer even very high yielding steps overall yields would suffer. In an effort to hold off this "wiring catastrophe" various manufacturers have begun to offer Cu rather than Al (2% Cu) metalization, and begun to explore "Low- κ " or low dielectric constant polymers and aerogels to push down the number of layers. However, this requires using nontraditional materials for both metals and dielectrics, along with a host of adhesion and diffusion layer barrier materials. In chip scale fabrication it is unclear whether this will lead to the desired improvements simultaneously in both yield and performance. The IBM 5HP process eventually provided 5 Al-Cu wiring layers on silicon dioxide, while the IBM 7HP provided 6 layer Cu process initially with Oxide, but eventually promises a "Low- κ " dielectric called SILK, for SiLow- κ .

During the course of this work, a key "second discovery" was made about the wiring in the Rockwell process. The dielectric used was a polyimide known as DuPont 2610/11. Early test structures fabricated with rather long wiring indicated that speed predictions for these lines were off by between 30 and 50%. The test structures were discussed in the predecessor final report. Eventually it was discovered that there were three problems. The first problem was that the polyimide was anisotropic. This meant that the dielectric constant measured horizontally in the plane of the polymer was higher since the polymer chemical chains lay horizontally in the film and electron shifts along these chains increased the dipole moment of the molecules in the film. The dielectric constant model that predicted speed to within 5% was one that assumed the vertical dielectric constant was 3.2 (DuPont claimed 2.7), and the horizontal one was 4.2, or larger than silicon dioxide at 3.8. This caused another retrenching of the design group very late in the design cycle. But in a sense discovery of these problems and publication of them is one of the missions of such research. A paper on this was published in TVLSI.

There are several strategies one can use to mitigate the effect of the "wiring catastrophe." The first is to try to "hide" wiring delays by "pipelining" them. That is during certain cycles an entire operation could consist of just wire transport for signals. Through the use of these pipelining cycles, "wire hiding" cycles can be conducted simultaneously with other more useful operations, so that only a portion of the pipelined operations are "wasted" just moving signals are around. However, inevitably this strategy exacerbates

the CPI penalty of branches that require “flushing” of many of these kinds of operations when they are unusable.

Another strategy is to push some of the wiring layers off the top of the on-chip interconnection stack into a package environment. Packages can provide a large number of “off-chip” wiring layers that “mate” onto the “on-chip” wiring stack, but can be fabricated independently and tested prior to this mating operation. The multi-chip package has provided a way to accomplish this. In fact, many of IBM’s mainframe computers made use of this by passing some of the longer wires thorough solder-bump connections into the wiring layers of the TCM, IBM’s ceramic multi-chip package. These package wires are fabricated on a much larger (mil) scale compared to chip-scale (micron) wiring dimensions. Up to 80 wiring and power distribution layers could be implemented in these packages. Of course, it is not necessary to package multiple chips for this strategy to work. Single chip packages can provide such supplemental wiring layers, provided that area array contacts are provides across the whole face of the chip.




Figure 33. Fast RISC Caching Strategy through L2 showing 1024-bit pair path between L1 and L2 for transferring a single line of cache between L1 and L2 in one cycle.

Because yield, speed and power limitations dictated such a small L1 cache a miss would be fairly likely, and hence to minimize the penalty of these misses a very wide transfer path to L2 was selected.

II.2.2. Technology

The technologies which are used in the F-RISC / G prototype, while providing the performance necessary to achieve its 1 ns cycle time and 2 GHz clock, also impose several difficulties in its design. Most notable among these limitations is poor device integration due to yield problems and comparatively high static power dissipation. However there were certain generic tools used, especially in circuit design, which worked and worked very well, and still remain today as viable approaches in SiGe HBT style design.

II.2.2.1. Current Mode Logic

F-RISC / G makes use of differential current tree logic and differential wiring published by Greub. These circuits are built out of differential pairs of transistors called *current steering switches*, which are arranged in a common emitter configuration.




Figure 34. Differential current switch.

Figure 34 shows a simple differential current switch. The current source I_s may either be passive (a resistor) or active (a transistor). By fixing this current, any inductance in the power and ground rails of the integrated circuit will see only this fixed current, and give no back EMF or simultaneous switching noise problems. Switching noise is becoming another terrible problem for conventional ON-OFF CMOS design.

F-RISC/G uses passive sourcing and a 0.25 V logic swing. Passive sourcing was selected due to the high V_{BE} of the devices supplied by Rockwell (1.35 V). Three levels of switches are stacked, allowing complex logic functions to be realized. It was desired that

F-RISC / G be compatible with standard ECL parts ($V_{EE} = 5.2$ V) so a passive source must be used.

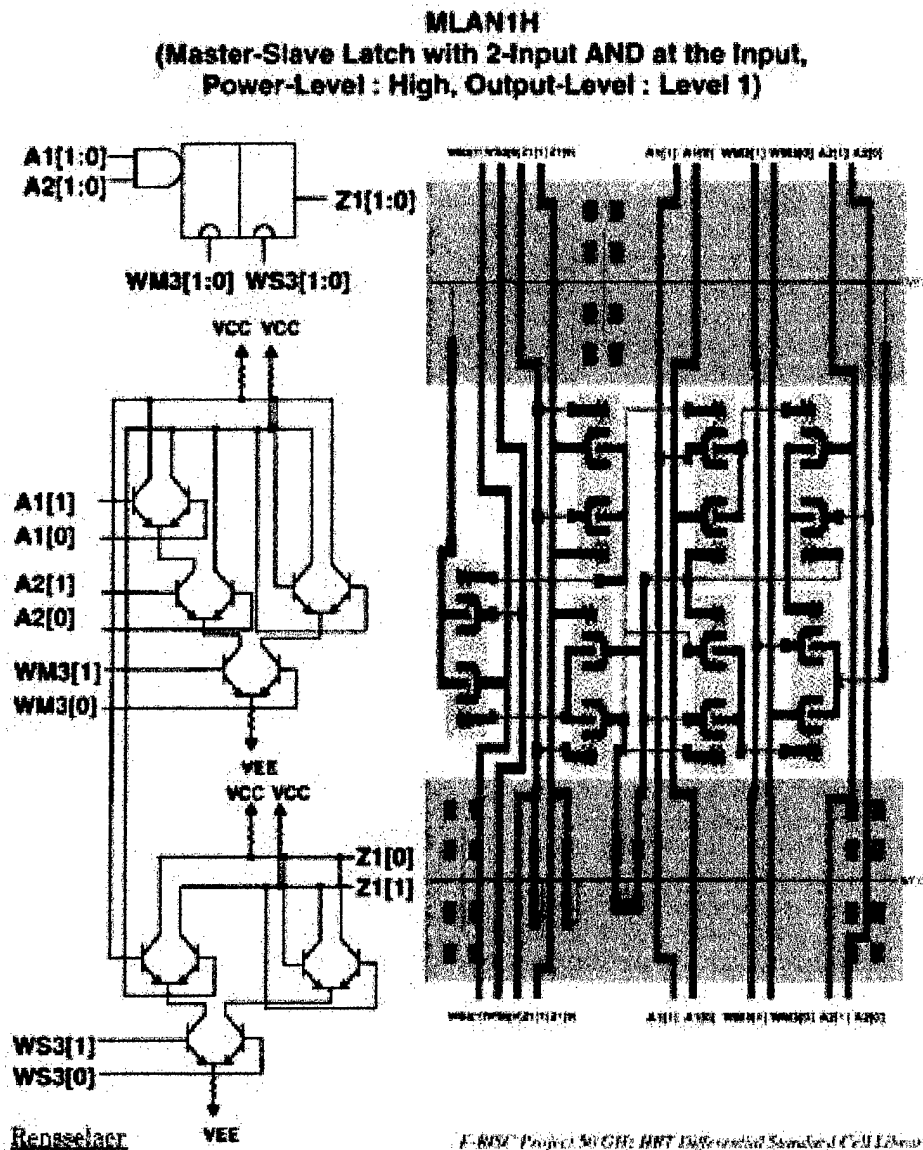


Figure 35. Sample Std. Height CML Cell from FRISC Cell Library.

Differential circuitry is used due to its common-mode noise immunity and the elimination of the common reference voltage required in Emitter-Coupled Logic (ECL). An added benefit of using differential wiring is that inversions can be accomplished merely by

flipping wires. The use of differential wiring can increase capacitance and requires more routing area. However, the utility of the full differential scheme is that it greatly reduces inter-wire coupling. In many cases HF HBT circuits built with single ended logic experience uncontrollable oscillation due to signal feedback in the high gain circuitry. With full differential wiring we have never had a problem with it. However these switches do have a hidden problem if there is skew between complementary pairs. A tiny amount of skew between complementary pairs can cause both transistors in the differential pair to cut off, causing an interruption of the current flowing reintroducing switching noise.. The high degree of functionality of the CML circuit class is exemplified in Figure 35 which shows a standard height cell from the FRISC cell library for the master-slave latch with two input AND gate at the Master input. Only two current trees were required to implement this storage element. The current fixing resistors and pull-up resistors near the power and ground rails set the current in the circuit, which is constant. The path through the tree for this fixed current is specified by which way the differential current switches steer this current through the tree. Only one path is active through the tree at any given time. The constant current flowing through the circuit that makes the circuit operation free of switching noise, is dependent on having very little skew between the differential pairs that enter the current tree at various levels. This requires routing of the differential pairs on paths that are essentially of equal capacitance. Note that the length of the lines is only one contributor to this capacitance, as the capacitance of each wire depends to some extent on the presence of other nearby conductors. However, to first order keeping the pair matched in length comes close to meeting this requirement. If this skew requirement is not met then current spike switching noise is reintroduced since the current switch can actually shut off all current if both signal and complement are simultaneously off. In addition, erratic behavior can be seen when clock and complement are simultaneously on.

This particular cell illustrates this exact requirement in the differential system. Specifically this cell has a pathological behavior in that if there is significant skew leading to the clock and its complement being asserted simultaneously, then the latch turns transparent. This behavior is also seen in standard latches, but is something that is often forgotten when working with full differential CML. The problem can also be introduced if rise and fall times for the clock and its complement significantly exceed the time for the latching action to take place. Rise and fall current paths often are dissimilar and so additional dissimilarities between complementary signal pair swings can arise. This results in clock and complemented clock appearing to be simultaneously in the forbidden region for switching for a period long enough for both signals to possibly appear asserted, resulting also in transparency. For this reason use of the MS latch incurs some risk that demands utmost attention be focused on never permitting significant skew or long rise or fall time to appear on the clock lines for this type of macro. Simulations of this type of failure become difficult since they must be for relatively large circuits, with realistic wire models, and done in SPICE to study the effect. This flaw in the behavior of the MS latch became especially important when it came time to test the chips. The MS latch was used in long shift registers of the FRISC "at speed" boundary scan test scheme. Rise and fall times on clocks for long shift registers arise in that circuit, which emerged as being extremely sensitive to this problem. This made testing of the chips extremely

difficult. Fortunately, a significant failure in the design came to light in this regard in the short period between when the glass plates for the fabrication were completed and when the wafer fabrication began. This involved the way in which the clock for the BS chain was routed in the chip, starting at the center and emanating to each side chain in the form of a T. At the edges of the T near the corners of the chip clock skew became significant, and it was actually possible to lose a bit along the chain at the corners. Fortunately one of the students in the group developed a clever way to fix the problem with just two new plates.

II.2.2.2. F-RISC / G Cache Implementation

While few of the design constraints on the F-RISC/G cache resulted from architectural issues, the *design* of the F-RISC / G core processor constrained the design of the cache to a great degree.

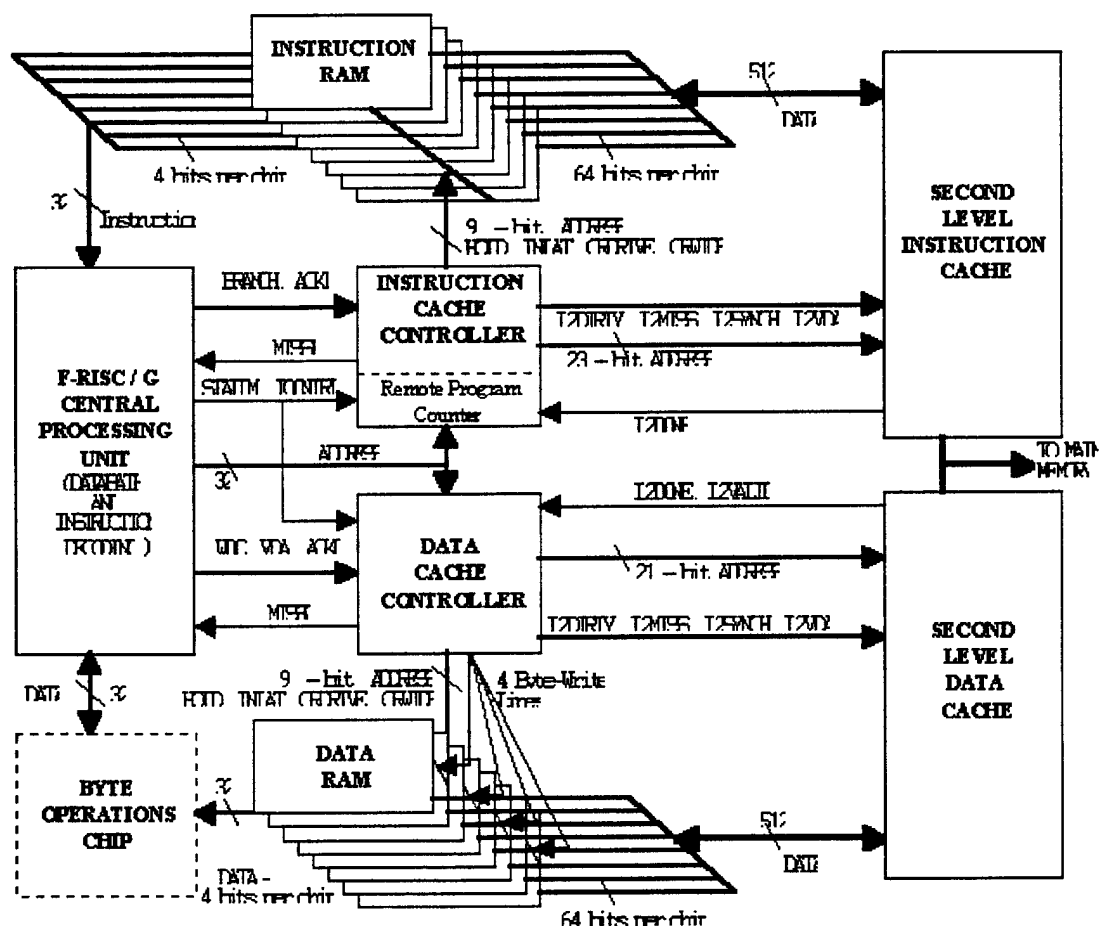


Figure 36. F-RISC / G System.

The F-RISC/G system is illustrated in Figure 36. The Central Processing Unit (CPU) is comprised of four datapath (DP) chips and a single instruction decoder (ID) chip. Instructions supplied by the instruction cache are decoded by the instruction decoder, which sends the decoded operands and control information to the datapath.

The data cache is used only for LOAD and STORE instructions (as with most RISC systems, F-RISC allows access to memory only through these instructions.). The Level 1 (L1) Cache is comprised of the primary instruction and data caches. Each cache consists of a single cache controller chip and eight RAM chips. Each of the two cache controllers must perform slightly different functions, but configuration circuitry is used to permit a single design to function in either the instruction or data cache. Each RAM chip is configured to store 32 rows of 64 bits and is single-ported. One unique feature of these chips, however, is that they have two distinct "personalities." Each RAM may read or write data four bits at a time using the DIN and DOUT buses. Each 64-bit row of memory may be filled one nibble at a time. A separate 64-bit bi-directional bus (L2BUS) allows reading or writing of an entire row at once. The wide bus is used to communicate directly with the secondary cache, and thus is less time critical than the four-bit bus that is used to communicate data directly to the CPU datapath.

II.2.3. Advanced MCM Packaging

Each cache must be able to handle one new memory access each cycle. Were the processor and cache to operate serially, this would require, for the data cache, that an address be communicated from the datapath to the data cache controller, that the tag be compared, that the address be forwarded to the cache RAMs, that the RAMs perform a read and multiplex the appropriate data to the output pads, and that the data be communicated back to the datapath in less than a nanosecond.

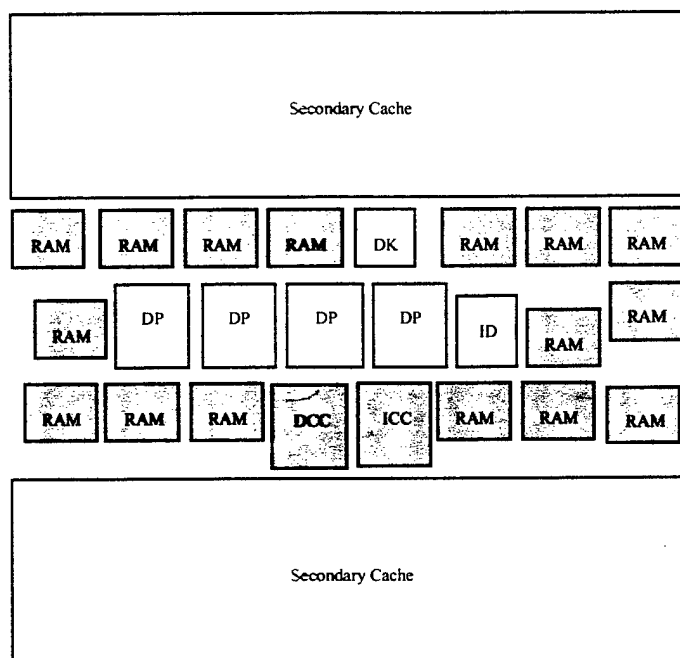


Figure 37. Possible MCM chip arrangement with L1 and L2 chips in place.

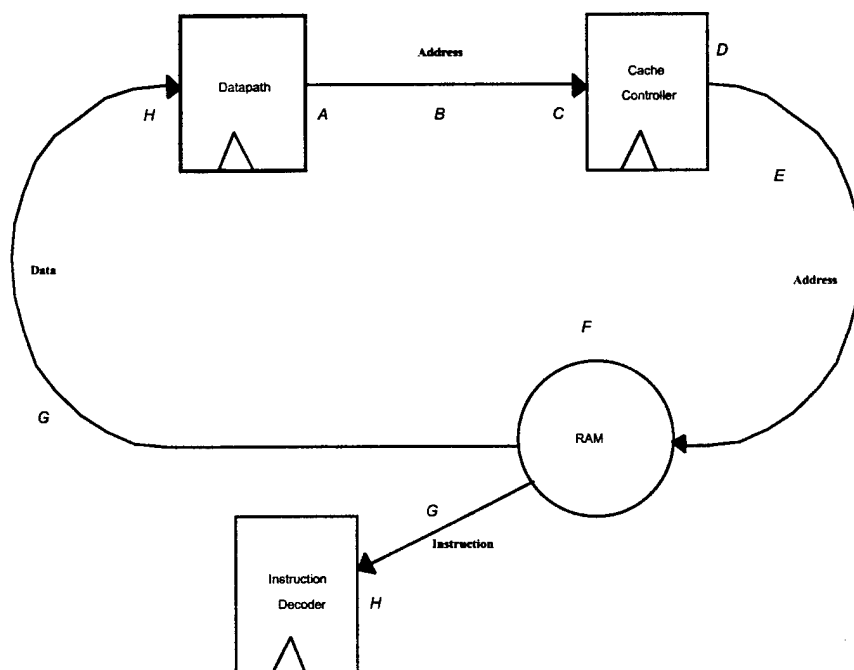


Figure 38. Critical path diagram.

Delay	Components of Delay
A	Driver Delay + On-Chip Skew
B	MCM Time of Flight + Skew
C	Receiver Delay + 2 Multiplexer Delays + D-Latch Delay + On-Chip Skew
D	Driver Delay + On-Chip Skew
E	MCM Time of Flight + Skew
F	RAM Read Access Time
G	MCM Time of Flight + Skew
H	Receiver + D-Latch Delay + On-Chip Skew

Table II.2.3-1. Delays along critical path.




Figure 39. Data cache critical path.

All of the memory subsystem data critical paths are shown in Figure 38 while this particular critical path is diagrammed in Figure 39.

The cache RAM blocks were designed to be accessed for read operation in 500 ps, and the cache RAM as a whole requires 750 ps from address presentation to data valid. This clearly makes it unlikely that the entire cache operation can be performed in 1 ns.

As a result, the cache and CPU are pipelined, so the effective allowed time for the data cache is 2250 ps (1850 ps-2100 ps for the instruction cache). Specifically, two CPU pipeline stages are allocated for each memory operation. The instruction fetch takes place during the **I1** and **I2** stages of the CPU pipeline. Data reads take place during the **D1** and **D2** stages, while data writes are additionally allotted the **DW** stage. The **D1** and **I1** CPU stages correspond to the **A** cache stage, while the **D2** and **I2** stages correspond to the **D** cache stage.

The data cache controller must be able to receive the address, latch it, run it through a multiplexer (which is used to select alternate address components in the event of a primary cache miss - specifically the tag stored in the tag RAM), and drive it onto the MCM lines. Allowing for slack and capacitive loading, 330 ps is a reasonable time allowance for these operations. A similar amount of time should be allotted to the datapath to drive the address and receive the data. This leaves approximately 840 ps for communications between chips. Note that the address transfer between the datapath and the cache controllers is further constrained by latch clocking to approximately 500 ps (or, more precisely, to approximately an integer number of clock phases - two phases seems to be the minimum attainable delay.)

Assuming a dielectric constant for Parylene of 2.65, polyimide, or BCB, the time of flight on the MCM would be 5.43 ps/mm . Allowing for clock skew between chips, rise time degradation of MCM signals, and some slack due to variations in MCM dielectric constant and dielectric thicknesses, an MCM time of flight of 5.75 ps/mm is reasonable for the purposes of this analysis. This would mean that the total MCM distance allowed for this critical path is approximately 146 mm. These times do not take into account the resistance of the lines which results in an R-C charging effect which increases rise time at both the drivers and the receivers; it is hoped that these lines will be wide enough to minimize this problem. If ρ is the interconnect metal resistivity, l is the line length, t is the interconnect thickness, and d is the dielectric thickness, the R-C charging effect can be approximated as:

Equation I.2.-23

$$\tau_{RC} \approx RC \approx \frac{\rho \cdot l^2 \epsilon_r \epsilon_0}{td^2}$$

Looking at this portion of the cache subsystem critical path more closely, the datapath chips and the cache controllers are each clocked by a global de-skewed system clock. The pipeline latch on the cache controller, which receives the address from the CPU, is clocked approximately 500 ps after the address is formed in the datapath. This means that there is 500 ps allowed for the datapath I/O drivers, the MCM time of flight, the cache I/O receivers, and associated skew, slack, and rise time degradation allowances.

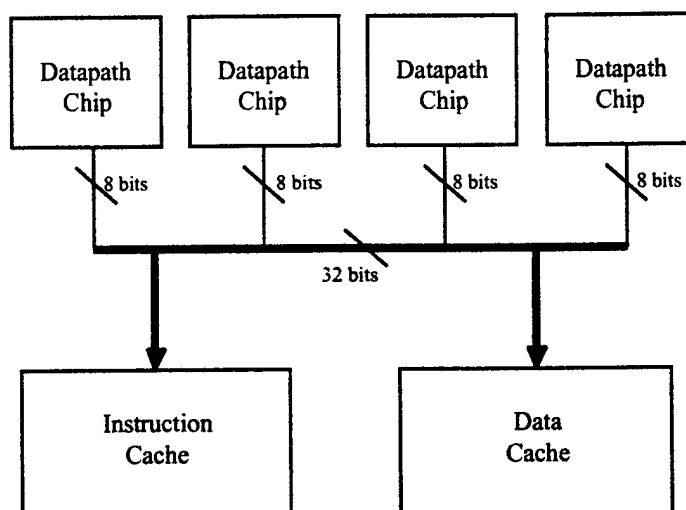


Figure 40. Address transfer from CPU to caches.

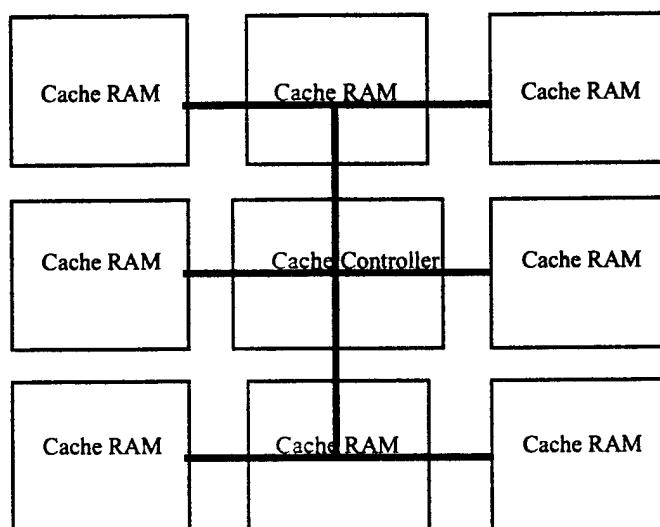


Figure 41. Single bus address transfer from controller to RAMs.

The next stage of the critical path is the transfer of the address from the cache controller to the RAMs. Each cache controller must send a 9-bit address to each of 8 RAMs. Were each cache controller to incorporate only one set of address output drivers, then this 9-bit bus must be long enough to reach each of the eight RAM chips, as shown in Figure 43.




Figure 42. Dual bus address transfer from controller to RAMs.

If the cache controller is given a second set of address drivers for this 9-bit bus, then the length of the longest address transfer from cache controller to most cache RAMs is significantly reduced (Figure 42) on the MCM. Also the parasitics associated with having multiple receivers on a given transmission line is reduced.

If a LOAD or an instruction fetch is taking place, then when the cache RAMs receive the address they are expected to read the appropriate location and send the data to either the instruction decoder (instruction cache) or the datapath chips (data cache). The CPU data and instruction word size is 32 bits, so in each cache each of the eight chips provides 4 bits of data.

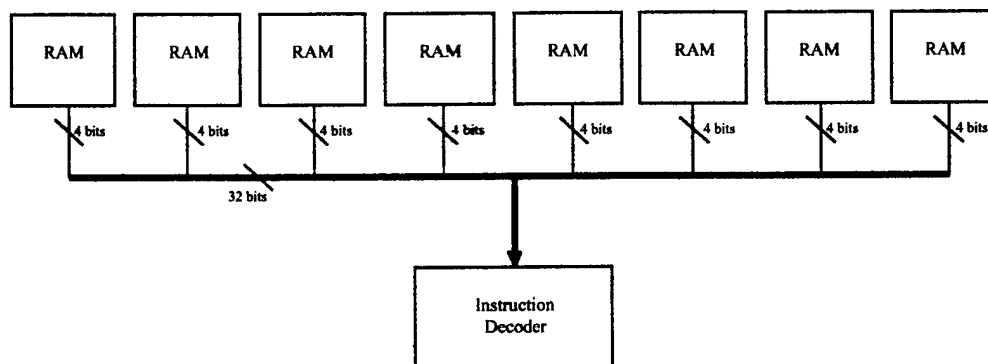


Figure 43. Instruction transfer - RAM to ID.

In the instruction cache, the eight cache RAMs must send four bits of data each to the instruction decoder (Figure 43). The length of the longest net for this portion of the critical path is determined by the longest distance between any RAM in the instruction cache and the instruction decoder.

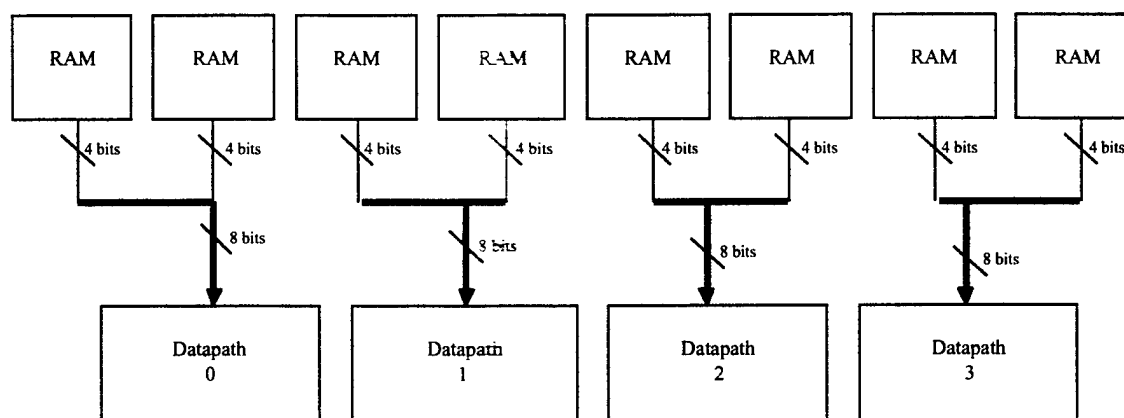


Figure 44. Instruction transfer - RAM to ID.

For the data cache, each datapath chip communicates with two data RAM chips. The length of the longest net for this portion of the critical path is therefore determined by the longest distance between a RAM in the data cache and its associated datapath slice. Since each of these nets must connect only three chips, as opposed to the instruction cache in which each net must connect nine chips, one would expect these nets to be shorter than in the instruction cache.

The constraints on the critical paths are

Instruction cache: (worst case)

$$1560 \geq D+E+F+G+H$$

Data cache:

$$1790 \geq D+E+F+G+H$$

Simulations based on preliminary MCM placement and routing predict a time of approximately 1584 ps for the data cache (including skew), which leaves approximately 206 ps for the byte operations chip should one eventually be incorporated. The predicted time for the instruction cache is 1504 ps on the fast path, and 1589 ps on the slow path (which has a constraint of 1675 ps). Table II.2.3-2 shows a breakdown of the timing for the cache subsystem critical paths.

		Data Cache	Instruction Cache	Instruction Cache
			Fast Bits	Slow Bits
A	Address I/O (datapath):	145	145	145
B	Address Transfer (DP to CC):	170	170	170
C, D	Address I/O (CC):	334	334	334
E	Cache RAM Address Transfer (CC to RAM):	300	300	300
F	RAM Access Time:	750	750	750
G	Data Transfer:	200	120	205
Total:		1899	1819	1904
Allotted:		2250	1850	2100

Table II.2.3-2 Critical path timings.

II.2.4. Cache Pipeline

The F-RISC/G CPU contains a seven-stage pipeline. Both the instruction and data caches are allotted two pipeline cycles to complete a fetch, and the data cache is allowed three cycles to complete a store. In the event of an acknowledged miss (a miss which is not ignored by the CPU due to an interrupt or trap) the CPU pipeline is stalled.

Table II.2.4-1 shows the operations, which take place in either cache during a fetch. Cache Controller and RAM operations may take place in parallel where appropriate.

Controller	RAM
Receive Address	
Tag RAM read	Receive Address RAM read
Tag compare	
Send miss	Send data
Wait for acknowledge	

Table II.2.4-1. Cache Operations During a Fetch.

The operations shown in Table II.2.4-1 can be divided into three stages as shown in Table II.2.4-2. Figure 28 shows cache operation over time if the cache is operating sequentially. The numbers in the table represent addresses sent by the CPU to the cache to be fetched. Although not every address will miss, it is assumed that the cache hardware and CPU / Cache interface require regularity of operations, so each address must pass through the miss handling stage. If each cache stage takes one cache cycle, then each fetch requires three cache cycles. In addition, the cache can only handle one address every three cycles.

By incorporating pipelining, however, it is possible to allow the cache to operate in parallel with the CPU. Although each cache fetch will still require three cache cycles, the cache can handle three addresses in any three-cycle period. By isolating the cache hardware through the use of "pipeline latches," it is possible to attain this type of behavior.

Stage	Controller	RAM
Read Address	Receive Address	
	Tag RAM read	Receive Address RAM read
Send Results	Tag compare	
	Send miss	Send data
Handle Miss	Wait for acknowledge	

Table II.2.4-2. Stages of cache operation.

	TIME →						
	1	2	3	4	5	6	7
Receiving Address	1			2			3
Sending Data		1			2		
Handling Miss			1			2	

Figure 45. Sequential cache operation.

Figure 45 shows how the pipelined cache would behave over several consecutive fetch requests.

	TIME →						
	1	2	3	4	5	6	7
Receiving Address	1	2	3	4	5	6	7
Sending Data		1	2	3	4	5	6
Handling Miss			1	2	3	4	5

Figure 46. Pipelined cache operation.

As can be seen from Figure 46, each cache stage is isolated so that at any given time it can deal with an address different from each of the other stages. While each address still requires three cycles, the cache is capable of completing a fetch during each cycle, under peak conditions.

II.2.4.1. Cache Hierarchy

Cache hierarchy has been extensively covered in earlier reports. F-RISC/G's L1 cache controller utilizes direct mapped cache strategy for simplicity and in recognition of the yield problems for GaAs/AlGaAs HBT's. This leads to Cache Controller (CC) chip sizes of about 17K transistors, around the outer limit currently compatible with yields of 20% at 10K transistors counts. After extensive simulations some of the conclusions for determining the other parameters of the L1-L2 cache interface are recapitulated here.




Figure 47. Simulations: 2 kB Harvard caches, direct-mapped, Spice trace.

For a given block size the best performance occurs when the bus width is equal to the block width. If the bus width were smaller than the block size, then multiple bus accesses, each incurring a miss penalty, would be necessary (unless a hardware-intensive buffering scheme were used - in which case occasional stalls would still occur.) Figure 48 is a graph showing the cache stall component of CPI as a function of block size and bus width for the Spice trace.

Figure 49, plotted on the same scale as Figure 48, shows that the magnitude of the CPI results obtained using the Tex trace is lower overall all than the results obtained with the Spice trace. Once again, at a given bus width, smaller block sizes seem to yield superior performance.




Figure 48. Simulations: 2 kB Harvard caches, direct-mapped, Tex trace.




Figure 49. Simulations: 2 kB Harvard caches, direct mapped, gcc trace.




Figure 50. Simulation results for benchmark suite.




Figure 51. 2 kB Harvard caches, direct-mapped, block size equals bus width.

The optimum point (1.41) occurs with a block size and bus width of 64 bytes. Given the estimated latency CPI component of 1.45, the total CPI taking into account both latency and stall CPI components would be 1.86. Figure 50 shows a plot of the weighted mean stall CPI for all three cache traces as a function of block size and bus width.

Having determined that the optimum configuration occurs when block size is equal to bus width, it is possible to plot the stall component CPI as a function of equal block sizes and bus widths as shown in Figure 52, where the results of each trace are plotted along with the weighted mean. From this plot it is clear that the minimum CPI occurs at a bus width and block size of 64 bytes (512 bits). It is also interesting to note that at half that size (32 bytes) the CPI is approximately 1.44, which is only 0.03 cycles per instruction worse than the optimal point.




Figure 52. CPI as a function of set size, block size.

Figure 52 is a plot of stall CPI as a function of set size and equal block and bus sizes for a Harvard architecture with dual 2 kB caches and copyback. From this plot it can be seen that larger set sizes tend to produce better CPIs, although the bus width and block size seem to have a larger effect on the overall CPI.

Figure 54 shows the effect of varying set size for the three block sizes and bus widths which provide the best results for 2 kB Harvard caches employing copyback and the timing constraints mentioned earlier. As can be seen, the CPI does not improve markedly as set size is increased beyond 4. The effect of moving from a direct-mapped (set size = 1) cache to a 4-way set associative cache, however, is fairly significant.




Figure 53. Effect of set size.

Figure 54 illustrates the effect of various cache architectures on stall CPI. The graph shows CPI as a function of block and bus width for a Harvard cache with 2 kB per cache, a unified cache with 4 kB of single-ported RAM, and a unified cache with 4 kB of dual-ported RAM. A direct-mapped cache employing copyback is assumed.

In a unified cache with dual ported RAM it is possible to read both an instruction and data simultaneously, while, for a single-ported RAM scheme, it is possible to perform only one access at time.

As the graph shows, the Harvard cache tends to perform the best. For the unified cache designs, the use of dual-ported RAM provides the best results except at extreme block sizes.

For the single ported unified cache, at most one cache access, instruction or data, can be accomplished at any time. As a result, the equation used to calculate CPI from the DineroIII output is as follows:




Figure 54. Effect of architecture on CPI.

DineroIII output is as follows:




Figure 55. Effect of architecture on CPI.




Figure 56. Effect of Harvard cache size on CPI.

Effect of Harvard cache size on CPI Figure 56 shows the effect of varying cache size given a Harvard direct mapped cache employing copyback and a 64-byte block and bus

width. The stall component of CPI drops below 1.5 at a cache size of 2048 bytes per cache. At twice that memory size there is comparatively little improvement in CPI, and there is little doubt that it would be extremely difficult to implement that much memory given the interconnect lengths that would be required and the difficulty in removing the heat from that many bipolar RAM blocks.

Based on these cache simulations, the design point which was chosen for the F-RISC / G prototype cache is as listed in Table II.2.4-3. Assuming a miss penalty of 5 cycles, the predicted stall CPI for this design is approximately 1.41.

Architecture:	Harvard
Ins. Cache Size:	2 kB
Data Cache Size:	2 kB
Write Policy:	Copyback, Write Allocate
Bus Width:	512 bits (64 bytes)
Block Width:	512 bits (64 bytes)

Table II.2.4-3. F-RISC / G primary cache parameters.

Table II.2.4-4 shows the results of the cache simulations broken down by type of event. The probability of each event occurring is also given. Based merely on these events, the stall CPI would add to 0.73. What remains unaccounted for are 68% of the instructions, which may be either ALU operations or BRANCHs. Each ALU or BRANCH operation can be assumed to take 1 cycle (since BRANCH misses are already accounted for in the "Instruction miss" category.) Therefore, the net stall CPI would be $0.73 + 0.68 = 1.41$, as reported above. Note that the "Reads" figure presented in Table II.2.4-4 includes the reading cycle that occurs at the beginning of each STORE, thus the write penalty would only be 1 additional cycle.

Event	# Occurrences	Probability	Cycles	Weight
Instructions:	2137409			
Reads:	440985	.21	1	.21
Writes:	254081	.12	1	.12
ALU / Branch:	1442343	.68	1	.68
Instruction misses:	77527	.04	5	.20
LOAD misses:	50635	.02	5	.10
STORE misses:	17881	.01	5	.05
Copybacks:	27179	.01	5	.05
TOTAL:		1.0		1.41

Table II.2.4-4. Stall CPI Components.

II.2.5. Model Deviations and their Implications

II.2.5.1. Re-Implementation with new models

As discussed in earlier reports for the F-RISC/G project the models first used to predict the performance of the processor had contained errors. These first surfaced when the results of the HSCD wafer fab were characterized. HSCD was a separate contract to Rockwell and Cadence to develop CAD tools for HBT's. As a result of this funding Rensselaer was engaged in designing several test chips. When the test chips returned from the foundry, circuits were found to run considerably slower than modeled. Before issuing the architecture chips for fabrication, the cause for this performance degradation in the test circuits had to be determined. Measurements obtained by the Mayo clinic and at RPI found the ILD thicknesses were thinner than stated in the design rule manual, causing an increase in the capacitance between the metal layers. The polyimide used for the dielectric was found to have anisotropic properties in the horizontal plane (at the temperatures used for deposition given the processing limitations imposed by the GaAs substrate). Interconnect capacitance was computed assuming a dielectric constant of 2.9 which was quoted by the foundry in the design rule manual. The dielectric constant in the horizontal plane however was almost 4.0 due to this anisotropy of the polyimide. The resistance of the metal interconnect was assumed to be negligible during the original implementation of the circuits. Neglecting resistance of the interconnect however added a significant amount of error to an extracted netlist. Finally, the devices were found to run

30 % slower than the foundry models specified. The combined effect of these discoveries so severely impacted the delay and skew in the processor, that it would not function at any speed. Analysis of the new critical paths given these effects showed that even if the signal hazards preventing proper operation were remedied, the maximum frequency of the processor would not exceed 200 MHz. This result would have been disastrous for the prime contract, which promised 1000 MIPS on a 2 GHz clock. It would have vitiated the entire argument in favor of pursuing HBTs as alternative devices in computer design.

This turn of events was accompanied by an improvement in the foundry interconnect process. The original process offered only two layers of Au interconnections. By the onset of the second major contract in the fourth year of the effort a third level had been added. Most analog circuits do not need more than about two levels of interconnections, but digital circuits in the LSI and VLSI size range need three or more. The introduction of the third level was a concession Rockwell made to their own increasing digital business flow. The interconnect widths and spacing of the metal layers could also be reduced, due to slight improvements in design rules. These processing advantages made it possible to salvage the processor and rework the logic.

The increase in capacitance within the horizontal plane caused by the anisotropic dielectric is more problematic in differential logic families (such as differential ECL and CML used in this processor) because the odd mode switching of the differential pair requires twice as much charge to move between the wires. Horizontal electric field lines between the differential pairs are increased relative to non-differential excitation, leading to greater sensitivity to dielectric anisotropy. This effect is referred to as dynamic capacitance because the switching appears as an increase in the capacitance between two differential pairs. A special version of the RLC QuickCAP tool developed by Yannick LeCoz was created to quantify the impact of the anisotropy. Other commercial tools for capacitance extraction did not even handle this case. With the use of this tool it became possible to predict observed circuit speeds with an accuracy of 2%. To remedy the problems with the dielectric anisotropy, a tool was developed to automatically shrink the interconnect wires in the processor to the new minimum widths. This effectively increased the space between the wires as shown in Figure 30, dropping the capacitance back to a more manageable level. This tool could not simply shrink every wire in the system however. The power and ground rails had to be preserved, and width had to be maintained on special high power gates to ensure we did not violate any current density design rules. The shrink tool had to be carefully crafted to preserve the widths of certain nets, while shrinking others. Furthermore, this tool had to parse through the entire hierarchy of the design, shrinking the appropriate interconnects throughout the underlying subcell levels.




Figure 57. Metal Geometries for Capacitance Compensation Strategy, due to Unexpected Anisotropy in Rockwell Polyimide (which Increases Lateral Capacitance).

Shrinking the wires to counteract the capacitance increase compounded the resistance problems in the interconnect. The impact of the increase in resistance of the Metal 2 and Metal3 interconnect layers were not significant. However, the increase in resistance in the thinner Metal 1 wires caused significant changes in the resistive aspects of the signal delay. A considerable amount of the rerouting involved replacing metal 1 routing with metal2 and the newer metal 3 layers.

II.2.5.2. QSIM Logic Simulator

The original VTI simulation tool used did not support distributed resistive effects in the interconnect, nor did it support signal skew along a wire with multiple receivers. Most of the intra-gate circuit resistance was assumed to be in the driver output, so the initial implementation ignored the metal interconnect resistance. The delay was then only a function of the output driver and the capacitive load it must drive. In the higher current high-speed HBT circuits, the resistive effects become a significant limiting factor in the operational speed of a circuit. To properly simulate the resistive effects in the interconnect, a new simulation tool had to be orientated for the HBT design process.

A second simulation tool (called QSIM) was available with the next version of our VTI CAD tool. The QSIM simulator specifies a net as either low or high, but does not indicate a transition phase. This had to be taken into account when computing the delay for nets that have a long signal rise time. Unlike the Mixed-Mode simulator used previously, QSIM netlists are extracted directly from the final layout, and may be simulated with the extracted resistances effecting circuit performance.




Figure 57. Metal Geometries for Capacitance Compensation Strategy, due to Unexpected Anisotropy in Rockwell Polyimide (which Increases Lateral Capacitance).

Shrinking the wires to counteract the capacitance increase compounded the resistance problems in the interconnect. The impact of the increase in resistance of the Metal 2 and Metal3 interconnect layers were not significant. However, the increase in resistance in the thinner Metal 1 wires caused significant changes in the resistive aspects of the signal delay. A considerable amount of the rerouting involved replacing metal 1 routing with metal2 and the newer metal 3 layers.

II.2.5.2. QSIM Logic Simulator

The original VTI simulation tool used did not support distributed resistive effects in the interconnect, nor did it support signal skew along a wire with multiple receivers. Most of the intra-gate circuit resistance was assumed to be in the driver output, so the initial implementation ignored the metal interconnect resistance. The delay was then only a function of the output driver and the capacitive load it must drive. In the higher current high-speed HBT circuits, the resistive effects become a significant limiting factor in the operational speed of a circuit. To properly simulate the resistive effects in the interconnect, a new simulation tool had to be orientated for the HBT design process.

A second simulation tool (called QSIM) was available with the next version of our VTI CAD tool. The QSIM simulator specifies a net as either low or high, but does not indicate a transition phase. This had to be taken into account when computing the delay for nets that have a long signal rise time. Unlike the Mixed-Mode simulator used previously, QSIM netlists are extracted directly from the final layout, and may be simulated with the extracted resistances effecting circuit performance.

Similar to the other tools used in this CAD package, the QSIM simulator is orientated for a low-speed CMOS process. Given a process technology file, the extractor extracts the resistance and capacitance for every net in the system. Upon simulation initialization, QSIM uses the process technology file to generate a delay database for every net in the system based on the interconnect information extracted, the driver width (this was developed for CMOS), and the total input capacitance of the input gate/s on the receiver transistors.

The first approach to orientating QSIM to our process involved generating a process technology file for a CMOS process that will behave like our HBT circuits under the operating conditions for our system. Do to the inherent differences between the two device families, it was determined that this modeling strategy would not yield accurate results.

Fortunately an option in the netlist description allows the delay from a driver to a specific receiver to be specified. This option is sometimes used to override the extracted delay for a CMOS simulation. The delay statement (as it will be referred to here) proved to be the perfect tool for mapping our technology to this simulator. The simulator delay calibration unit was disabled, and all of the delays were manually inserted directly into the netlist. These delays were calibrated by running Spice on every driver to receiver pair in the entire netlist. A tool was developed to take the extracted resistance, capacitance, and driver circuit for every receiver in the chip and run SPICE on them to determine the worse case propagation delay from driver to receiver. This delay was then inserted into the netlist, and the logic vectors could then be run to determine if the system still worked with the new extracted RC delays.

Interconnect resistance is extracted from the layout by adding up the all of the via resistances for every via the signal must traverse, and adding it to the metal interconnect resistance. The interconnected resistance of each metal layer is determined by measuring the total length of metal the signal travels through, dividing it by the width of the metal, and multiplying by the interconnect sheet resistance. Three dimensional capacitance extraction using QuickCap was used to tune the capacitance extraction tool used in the Compass tools for the thinner wires used to counteract the capacitance increase of the anisotropic dielectric. Extracting the capacitance of the entire chip using the QuickCap capacitance extraction tool requires a week of computation on our fastest workstation. Since the design methodology involves several hours of simulation, followed by logic and clock optimization, and then a re-extraction of the chip, waiting a week for accurate capacitance extractions was not feasible. Instead, the fast though less accurate extraction using the Sakurii theorem permitted reasonably accurate extractions in less then an hour, which worked well with our design methodology. The capacitance of each signal was determined within the horizontal plane for each layer of metal. The capacitance numbers in the horizontal plane were then offset by a correction factor that assumes that one of the signal's neighbors is swinging in the opposite direction (the other wire in the differential pair) and the other neighbor is constant. This is necessary because the simulated delays are static, and are not dependent upon the transition of the neighboring signals. Next, the capacitance between the interconnect layers was extracted and added to the horizontal capacitance results for each signal. Though the worse case signal delay occurs when a

signal swings in the opposite direction as both of its neighbors, modeling the interconnect delays with only a single opposing neighbor provides accurate results. The interconnect was redesigned so two signals undergoing state transition at the same time were not routed next to each other. The separation of these signals was not difficult because signal skew and the fact that there are four different clock phases upon which signals undergo transition, made the separation of such signals easier. Cross-talk in future designs implementing a single phase clocking scheme may produce a more hazardous effect in the system increasing the complexity of the design implementation. Figure 58 shows a second method of reducing cross-talk in routed differential pairs.




Figure 58. Reducing Signal Cross-Talk.

A differential pair is swapped to reduce the cross-talk between neighboring signals by balancing the dynamic capacitance between the two signals.

If too many signals switching at the same time occupy a dense routing region, it may not be possible to separate simultaneously switching signals. Under these circumstances, a crossover may be used to swap which differential pair neighbors the transient signal, thereby reducing any cross-talk between the signals, which may occur.

II.2.5.3. Top Level Full Speed Simulation

Once the tools were calibrated and the interconnect re-extracted, it was found that the delays on many nets increased several hundred percent. The entire core of the processor had to be re-implemented, and the critical logic paths had to be optimized using new circuit techniques. Verifying the new design iterations initially began using the single chip test vectors designed to determine correct functionality of the original chips (keeping in mind the errors that these vectors were found to miss). Once the initial re-placement and re-routing was complete for the ID and DP, and the initial critical paths were addressed, the design was immediately moved to a top level simulation.

The top level netlist was developed to simulate the ID and four DP slices in a single system with the MCM delays included. The top level netlist was extracted from the schematic generation tool used to describe the top level interconnect. A "dummy" chip for the ID and four DP chips are included in this top level schematic, and provide connector information for the underlying chips. A single level hierarchical netlist is then extracted composed of only the top level interconnect and subcell calls to each chip.

Connector information is added to the netlist of the chips extracted from the layout. Since the layout of each chip was considered the "top" level of that chip, previous connector information was not extracted. Each pad must have a connector added according to the interconnect name used in the top level schematic. The top level MCM netlist must then have the corresponding subcell calls altered so it calls the extracted file types from an extracted layout instead of a schematic netlist (it expects schematic netlists since it itself was generated from a schematic). The netlist must then be flattened effectively generating what is essentially a layout extracted from the MCM down to the GaAs/AlGaAs circuits. The delay information for the nets within the core chips must then be re-inserted into the flattened netlist since the netlist flattening routine discards this information (a bug in the netlist utility of the CAD software). Finally the top level MCM delays are inserted into the netlist and top level simulation may be performed.

Unlike previous chip level simulations, the top level simulation provided a simpler external interface to the circuits being tested. In the top level simulation, the complete instruction (as it would be assembled) is applied to the Instruction Bus of the instruction decoder chip. This effectively inserts the instruction into the pipeline at the beginning of the Decode stage. Using this testing method, simple test vectors which a user can easily relate to are used, as opposed to the complex signals of the internal processor necessary to test a piece of a bit slice design.

II.2.5.4. Logic Path Optimization

The drastic degradation in performance of both the interconnect and the device made it necessary to completely re-implement the processor's logic and clock distribution. Furthermore, several logic errors made in the original implementation had to be corrected. Clock distribution both on-chip and across chip became a greater concern since there was not sufficient data to determine the extent of the process variances.

Communication between two chips with a master clock whose phase error is not within tolerance (30 ps) became a concern since the clock deskew implementation made optimistic assumptions in the original design.

Using the new circuit extraction techniques, the design began with an analysis of the current on-chip critical paths. This includes logic paths contained entirely in the instruction decoder or datapath chip. This analysis ignored clock distribution and measurements were made on which logic paths did not make speed. The results of this analysis showed that of the paths which were extremely slow, up to 80 % of the delay was contained in the interconnect.

Circuit analysis shows that the resistive effects of the interconnect created a significant amount of the delay in the processor. Most of this resistance was found to be contained in the lowest metal layer since this layer contains the highest sheet resistance in the process interconnect layers.




Figure 59. Solving Critical Paths with Pipeline Adjustments.

Logic paths that do not reach the destination pipeline latch in time may be moved to the following pipeline stage provided the input remains stable and the output is not needed immediately.

The chips had to be rerouted, optimizing inefficient routing traces, and converting Metal 1 wires to Metal 2 or Metal 3 wherever possible.

Most logic paths in the system were still too slow to meet the target clocked 2 GHz processor cycle time. Several techniques were implemented to reduce logic delays along certain paths. Figure 32 shows a technique where logic is moved to the next pipeline stage in the processor. This method works provided the result for the pipeline latch is not used for other logic in the processor. Furthermore, the data on the critical net must remain stable during the following pipeline stage. If this is not the case, it is necessary to latch the critical net for the evaluation in the following cycle, adding an additional gate to the system.




Figure 60. Splitting A Master-Slave Latch to Alleviate Time Constraints. A master-slave latch may be split to provide a slow logic path with more time to the input of the slave latch.

Evaluating segments of logic a full pipeline cycle later than originally designed is difficult to do since the evaluated logic is often required that cycle. A second technique, which is more easily used in a design, is shown in Figure 61. In this case, a master-slave latch is split, and the logic evaluating the critical net is inserted between the master and slave latch in the circuit. With the use of a 4-phase clock, master-slave latches may be clocked on separate clock phases resulting in a system where coherent pipeline stages starts to break down. The system becomes a micro-pipelined set of latches tuned to providing the necessary time for the critical paths to properly evaluate. A similar

technique called "cycle stealing" is described in the following section on clock distribution.

Figure 31 in section II.2.1 shows a block diagram of the processor, outlining the new critical paths. The control signals are typically high fan-out nets and the capacitive impact of the wires significantly decreased performance of these signals. Other critical paths included the ALU result feed-forward path from the EX stage. The ALU result must be calculated and the result must be routed back to the input latches of the ALU for the next operation. These critical paths will become augmented once the clocks are re-distributed and the cross-chip clock skew is taken into account.

Clock skews between chips (described in the next section) complicate logic paths that make chip crossings. If the clock driving the signal is out of phase and behind the chip receiving the signal, the latch may be closed before the data is stable at the input, resulting of a loss of the information. Likewise, if the driving chip is ahead of the receiving chip, hold-time violations may occur, and the data may be corrupted in the receiving latch. To avoid problems with clock skew, the window when valid data may arrive was increased to 125 ps (from 30 ps). As much as 125 ps had to be removed from the time available in the logic paths, which involve MCM crossings. Logic paths that were not considered critical paths suddenly became far too slow to meet the target cycle time.

II.2.5.5. Distribution

The slowest logic path between two latches determines the limitations on a processor's speed. The clock skew and jitter (uncertainty in the master clock frequency from cycle to cycle injected by the external clock) significantly complicates this system limitation. Jitter can cause problems in systems when the latch setup or hold time is violated along a path that is tuned to barely make the target speed. Problems involving clock skew may be circumvented to a degree by carefully analyzing the data arrival times to each latch and balancing the clock edges with respect to these data arrival times. Process variations and thermal variations however, will cause additional clock distribution problems. These effects may alter the data and clock propagation time to the latches causing setup or hold time violations despite careful effort to analyze and design for the circuit's various parasitic components.

In a multi-chip system, problems with clock skew become even worse. Distributing the clock to all chips in an even manner is impossible when process variations are accounted for in the MCM technology. To account for this, an active clock deskew method is necessary to compensate for thermal and processing variances from chip to chip and across the MCM.

The clock deskew method proposed in the FRISC/G architecture uses a multi-channel delay locked loop to distribute the clock in a low skew manner. This Deskew chip contains a separate clock channel for each die that accesses the master clock (ID, 4 DP, ICC, DCC). The clock is sent to each chip and is returned on the return path for each die.

The delay is measured, and the on chip delay locked loops track the phases so each phase arrives at each chip at approximately the same time.

A SYNC signal is used to keep the on-chip system clock locked in phase 1 while the deskew chip's phase lock stabilizes. Once the on-chip master clock is deskewed, the SYNC signal is disserted, allowing the processor's four-phase system clock to activate.

Spice simulations of the Deskew circuit indicate that the clocks will arrive at the 4-phase generator of each chip in the circuit within 5 ps of each other [NAH 93]. These simulations however are very idealistic and don't account for possible thermal and process variations within the Deskew chip itself. However, this Deskew method will track and correct potential variations along the different clock paths external to the Deskew chip. The design of the active clock deskew circuit is a separate focus of research which to date has not been completed. To account for the lack of an active deskew circuit, work was done to implement a passive clock distribution method on the MCM to minimize the cross chip clock skew. Currently we have no resultant foundry data of our own on the expected variation of the MCM dielectric. The core processor was designed to tolerate a 125 ps cross-chip clock skew, which is 12.5% of the processor cycle time. The original design implementation could only tolerate 30 ps of chip to chip clock skew [Phil93].

To alleviate problems with clock skew between chips, additional slack in the chip crossings were added. This slack provides extra time during cross-chip transactions. Should the signal depart later then expected, there is still enough time to reach the destination chip before the data is latched. Furthermore, should the data leave the chip sooner then expected, it will be latched on the destination chip before the data changes in the next cycle.

II.2.5.6. Clock Distribution Strategies

The large clock skew between chips in a multi-chip system is somewhat counter-balanced by the long data delays from chip to chip. When the data propagation is estimated to be large, this situation is designed for, and a variation in the clock skew is less likely to result in setup or hold time violations. On-chip clock skew, although significantly less then the chip-to-chip clock skew may also result in setup or hold-time violations. Two strategies to distributing the clock to the system are described below. The first method distributes the clock in an H-tree shown in Figure 62. The H tree depends largely on symmetry of the system, and can be costly in wiring resources, particularly in a multiphase clocking system. In this technique, the clock comes in off-chip and is buffered. It is then balanced and sent to the ends of the chip. Active circuits may also be used in this scheme to match the on-chip clock with the in-coming master clock.




Figure 61. H-tree Clock Distribution. Local clock buffers are driven by a system clock buffer through a balanced H-tree of interconnect. Clock distribution is given precedence over routing resources to insure a balanced clock network.

A buffer tree may also be used, which provides more control of the clock skew in the local system, but relies more heavily on the CAD tools to correctly model the parasitic effects. The clock buffer strategy is used in the FRISC/G processor as well as the DEC Alpha. The Alpha chip used a distribution of clock buffers to feed a uniformly deskewed clock to each side of the chip [Gron98] whereas the FRISC/G processor routes buffered clocks based on careful timing analysis.

The clock must be distributed locally as well. Running the clock in the opposite direction as the data helps reduce problems with transparent latches, however, running the clocks in the same direction as the data allows for a faster implementation. Figure 62 shows the difference between the two.




Figure 62. Local Clock Distribution. Aggressive clock distribution may increase performance by skewing the clock to match logic propagation delay, however this makes the system sensitive to hold time violations.

Assume the Latch propagation delay is 30 ps, and the setup time is 35 ps. The slowest path of logic from latch to latch determines the limitation on the clock speed for the case where there is no delay on the clock between the latches. The required period of the clock (T_{clk}) for the logic shown in Figure 31 is determined by the following equation:

$$T_{clk} = T_{platch} + T_{slatch} + T_{plogic} \quad (1)$$

Where T_{platch} is the propagation time through the latch, T_{slatch} is the setup time of the following latch, and T_{plogic} is the propagation time through the slowest branch of logic. For this case, T_{clk} would be 185 ps.

The presence of clock skew alters this analysis. If the clock propagates in the opposite direction as the data, the equation is as follows:

$$T_{clk} = T_{platch} + T_{slatch} + T_{plogic} + T_{skew} \quad (2)$$

Resulting in a 200 ps T_{clk} .

Aggressive clock distribution however may use skew to actually increase system performance. If the clock propagates in the same direction as the data:

$$T_{clk} = T_{latch} + T_{slatch} + T_{plogic} - T_{skew} \quad (3)$$

Now T_{clk} is only 170 ps, 15 ps faster than the system with no skew at all on the clocks. Care must be taken when using this technique however. If the skew on the clock is large enough, and if the propagation through the logic is fast enough, this can lead to hold time violations, and the data in the latch could be lost. A transparent latch is the degenerate case, where the new data from the previous latch arrives at the latch input before the clock, thus eliminating the logic result before it is written to the latch.

The FRISC/G processor uses a four-phase system clock to provide control for a single port register file. The four-phase clocking scheme however, also provides greater control over the distribution of clock edges. In a four-phase system, master-slave latches are clocked on odd or even clock edges reducing potential problems with setup and hold time violations.

The clocks were balanced to maximize the setup and hold time window. Figure 37 shows the safest time to change the data on a falling edge triggered flip/flop or positive level sensitive latch.




Figure 63. Clock Transition with respect to data. To minimize possible setup or hold time violations, the clock must arrive in the center of the "safe transition" window.

The data transition must occur after the hold time from the previous falling edge clock cycle or the information will be lost. Likewise, it must be stable prior to the latch setup time before next clock cycle's falling edge. The processor's logic and clock delays were

analyzed to determine when the data would become stable, and a clock was routed with the proper skew such that the data would undergo transition at the midpoint between the setup and the hold time of the latch. Using this technique of optimizing signal and data transitions, a system can be developed to tolerate the maximum amount of error in timing analysis. This is an important aspect of the processor since previous experience has shown that the foundry process varies from the expected simulation models considerably. The modeling of the process has therefore been developed here (as opposed to using the incorrect models provided for us) based on limited foundry results (two wafers from a single lot). Hence, all estimates regarding process speed has been made conservatively in the event that an unforeseen effect causes the system to slow down.

II.2.6. Fabrication of the Chips

Ultimately after 3 years of model revisions, re-verifications, and redesigns of the 4 architecture chips a final reticle was prepared. Tape out occurred in January of 1998. However, even this reticle was not acceptable. Rockwell revealed to us last minute yield updates that suggested that a switchover to 2 micron by 2 micron emitter stripes would be necessary to assure yield for our rather large HBT count circuits. This required additional rework through out the spring and following summer

In late August of 1998, after one more last-minute discovery of a flaw in the skew of the boundary scan circuit clock, the final reticle was released to production. Such last minute discoveries can be expensive. In this case one of the students found a two-plate workaround that cost only \$4000 to correct.

Chip	Height(mm)	Width(mm)	Area (mm ²)
Datapath	9.457	8.457	79.978
Cache RAM	9.347	6.703	62.653
Cache Controller	8.365	9.472	79.233
Instruction Decoder	8.742	7.672	67.069

Table II.2.6-1 Dimensions of each of the F-RISC/G chips.




Figure 64. Final F-RISC/G Fabricated Reticle, showing the four main architecture chip sites. Right strip of smaller chips contains another copy of the RPI Test Chip and deskew test chip.




Figure 65. The Data Path (DP) Chip.

The Data Path chip or DP is the most sophisticated part of the architecture containing the adder and shift circuitry as well as the register file. Due to yield limitations the data path is organized in a byte-slice fashion. Each of the 4 chips in the 32-bit architecture captures 8 bits of the width of the processor. Hence the register file shown in the upper right of midline of the DP is only a 32w x 8b file. The artwork for the file was taken from the last HSCD retuned shared reticle fab, and was certified as a 5 GHz file. The adder had also been validated at 750 ps, of 666 MHz through a carry chain ring oscillator. One can observe in the component overlay diagram in Figure 66, that a fair percentage (about 20%) of the roughly 10,000 HBT devices used in this chip are for boundary scan

testing, scan buffers, and the four phase clock and clock distribution. Each chip had approximately this same overhead devoted to testing. The system architects for F-RISC/G worked diligently to anticipate the difficulty of finally testing and verifying the chips after fabrication. Experience in the earlier test structure fabrication experiments helped greatly to heighten awareness of the problems faced once the chips are fabricated including guaranteeing the ability to probe key signals with probes that can actually be purchased, and which work at the speeds required. Another consideration was the inclusion of sufficient pads that could be probed at DC to supply enough power and ground connections that voltage droop on internal chip power rails would not be severe. Special outrigger pads for probes were introduced for probe touchdown since this scratches the surface of the metalization. Inbound pads were included for MCM via attachment.

	Cache Controller	Cache RAM	Instruction Decoder	Datapath
Devices	10572	9679	7358	9785
Power (mW)	12633	12179	11573	12798
Area (mm ²)	79.23	62.65	67.07	79.98

Table II.2.6-2 Power Consumed & Area of each of the F-RISC/G chips.




Figure 66. Component Overlay for Data Path (DP) Chip.



Figure 67. Zoom on Figure 66 showing close-up photomicrograph of 32w x 8b register file, register file decoder on left and some of the Boundary Scan support circuitry, towards the left of the file.

Ultimately, the key test used to validate the timing and functionality for F-RISC/G DP chip consists of locking in the control signals of the instruction decoder for an ADDI immediate instruction with a fixed constant loaded in the immediate register. The add instruction adds the constant to the register file, and deposits the sum back into the register file. By monitoring the carry out signal for a specific constant, a periodic carry output signal is observed that can be validated.




Figure 68. Instruction Decoder (ID) chip.

Figure 68 shows the fabricated Instruction Decoder chip. The schematic for the instruction decoder is shown in Figure 69. It can be seen that the decoder is organized as a pipeline structure arranged around the 7-stage pipe structure discussed earlier. The actual decoding of control signals for each instruction pipe stage consists of only a few current trees after each pipeline flip flop plus wire loading. At 25 ps per current tree decoded lines are available about 75 ps after each clock phase. Key testing of the ID consists at speed of checking the FSM and four-phase clock generator, along with static decoder tests at DC. Fortunately, despite the pipelined structure of the F-RISC/G it was designed such that there would be no lower limit on clock frequency, and the architecture could be single stepped at essentially DC. No time of flight propagation delays were

relied upon for critical timing. This meant that provided the four-phase generator met speed the rest of the chip could be tested at DC for functional correctness in the full knowledge that speed would be met by the shallow decoders. In other words this part of the design was created with a great deal of slack for ease of testing.




Figure 69. Schematic of Instruction Decoder Chip.




Figure 70. L1 Cache Controller (CC) chip. Identical chip used with personalization for Data and Instruction Cache.

The cache controller chip is designed to manage address decoding for the data and instruction cache L1 chips. The cache controller must intercept addresses and identify whether the corresponding data or instruction items are in one of the L1 cache lines, and if not to cause appropriate action to take place initiating a full line transfer (in one L2 cycle time) of 32 words simultaneously. Each chip has a front door facing the ALU and a back door facing L2. A miss causes one of these transfers to take place by initiating access at the target address in L2. All 32 words are then transferred while the ALU must stall. The impact of this has been discussed earlier.


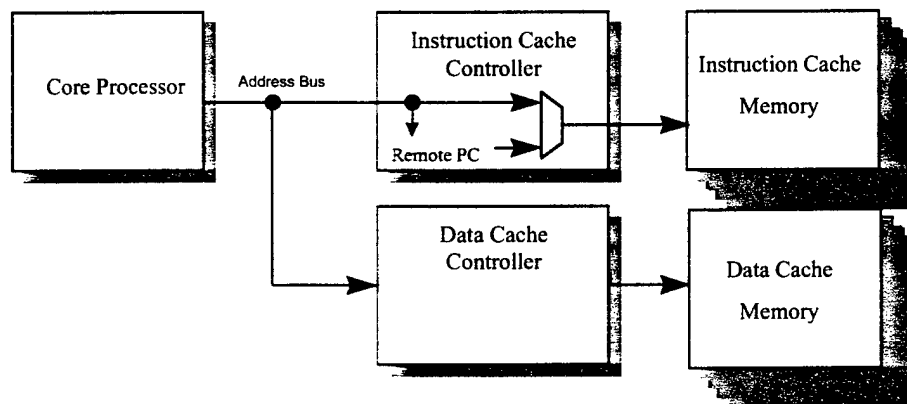


Figure 71. Zoom for Figure 70 in the vicinity of the Tag RAM files.

	Devices	Power (mW)
I / O	2548	2810
Write byte decoding	80	85
Tag RAM blocks	3420	4000
Testing logic	1068	1069
Control	410	467
Pipeline and RPC	2664	2208
Clock distribution	78	1714
Comparator	304	280
TOTAL	10572	12633

Table II.2.6-3. Cache controller device count comparison of F-RISC / G chips.



The cache controller was designed for use in both the instruction and data caches. For this reason the first pipeline latch serves also as the Remote Program Counter (RPC) in the ICC configuration. Figure 72 shows the manner in which the two caches share a common CPU address bus and how the RPC can be loaded from this bus. If two separate cache controller chips had been designed, it would have been possible to include only two pipeline latches in the DCC as at any given time only two addresses need be stored (the third always being available on the bus.) Since the hardware for the RPC had to be included, however, it was decided that it also act as a latch in order to reduce problems caused by hazards and skew on signal lines while at the same time minimizing chip configuration and initialization logic.

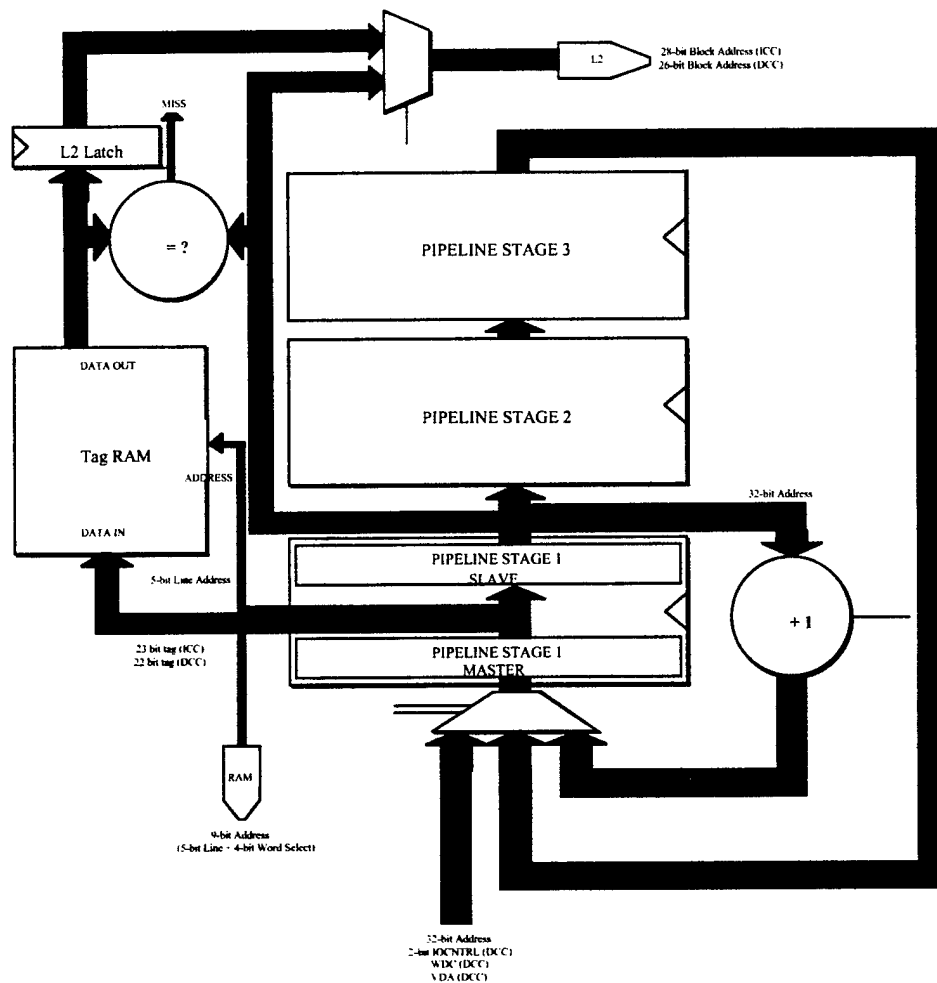


Figure 72. Simplified cache controller block diagram.

Although the responsibilities of the two cache controllers differ slightly, it was decided to design a single, configurable controller, due both to the cost and time required to design an extra chip; the operation of the controllers in both caches is similar enough that methods were found to minimize the penalty for using a single chip.

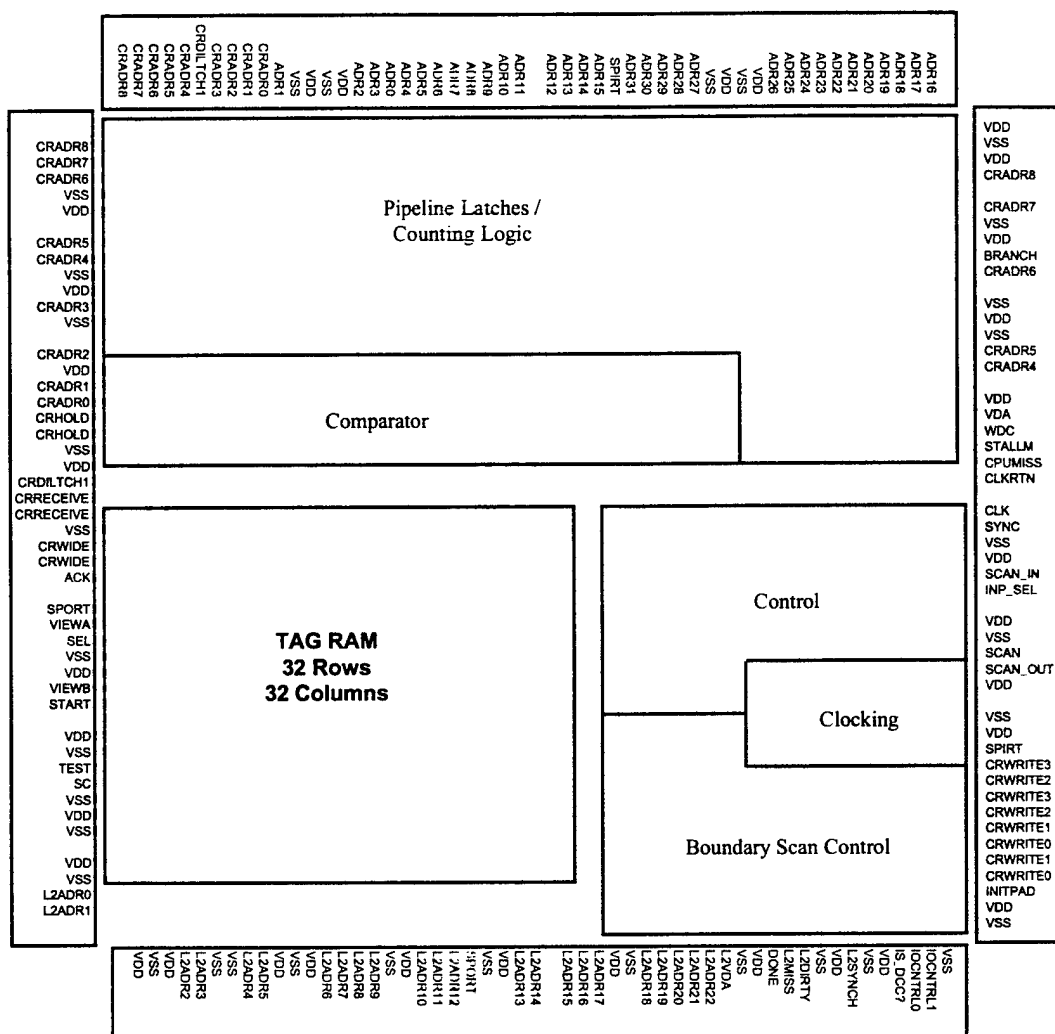


Figure 73. Floor Plan for Cache Controller Chip.




Figure 74. Universal L1 Data and Instruction Cache Chip.




Figure 75. Cache RAM chip Layout.

Unlike most boundary-scan schemes, the sampling and input latches are located in the core rather than in the pad ring. These latches and associated multiplexers and control circuitry take up most of the standard cell area.

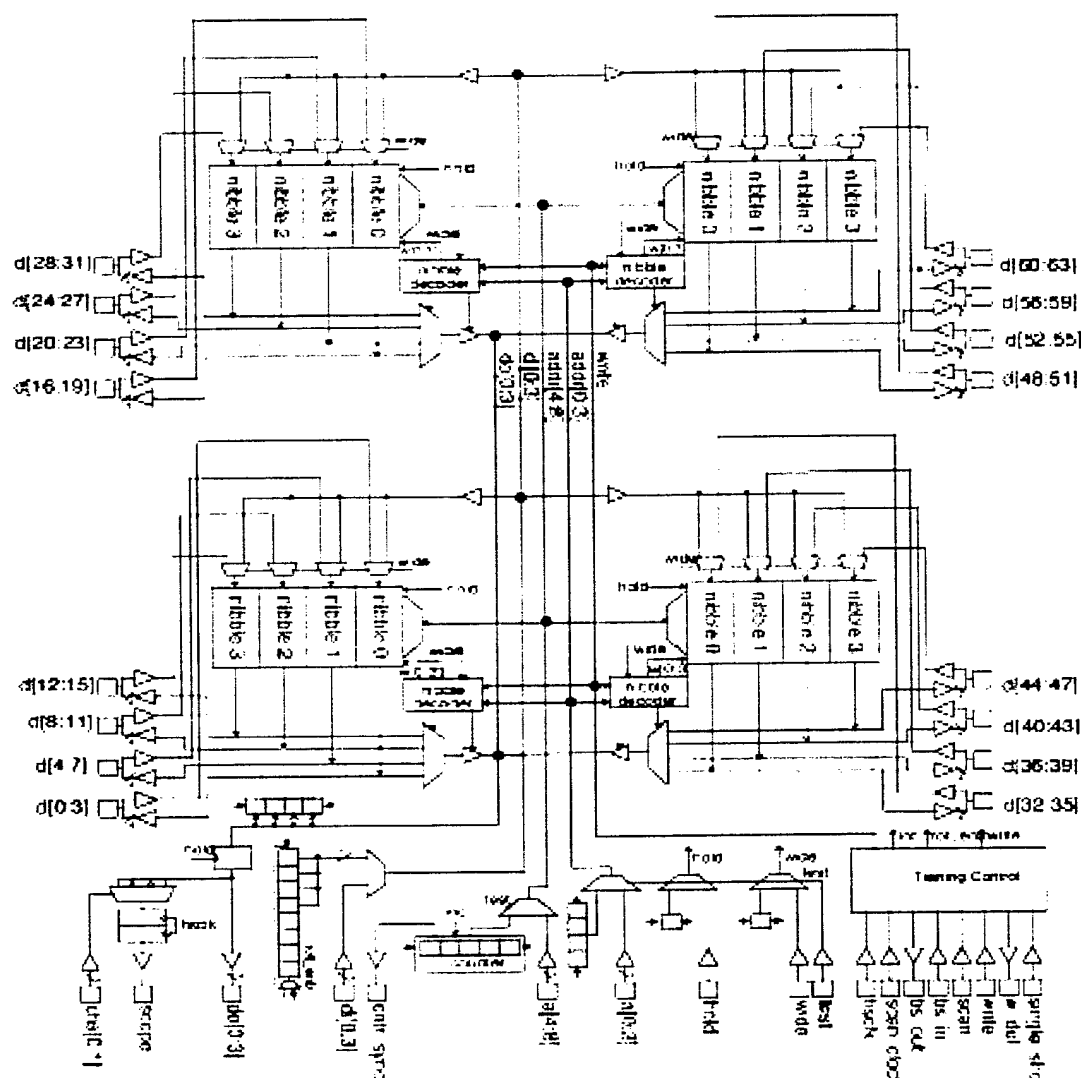


Figure 76. Schematic Diagram of Cache Memory.

II.2.7. Instruction and Data Configuration

The cache controller contains a pad, `IS_DCC?`, which is used to enable the chip to be configured for either the instruction or data cache controller. For data cache use the signal is asserted by hardwiring it on the MCM.

Additionally, when the chip is intended for the data cache, the `BRANCH` pad should be asserted by hardwiring it on the MCM: the ICC will have the `BRANCH` signal asserted by

the instruction decoder whenever a branch is to occur. This signal is used to determine whether the first pipeline stage (the remote program counter) is loaded or counts.

Since it is impossible to perform a STORE into the instruction cache, the WDC line must be hardwired low. In addition, the instruction cache must retrieve an address on every cycle, so VDA should be tied high.

II.2.7.1. Cache Controller Design

The cache controller chip is 8.365 mm high and 9.472 mm wide. Table II.2.6.-3 shows an approximate device usage breakdown for the cache controller chip.⁴ As in the cache RAM, a large percentage of the power is dissipated in the RAM blocks and the I/O pads.

Table II.2.6-3 compares the critical features of the F-RISC / G chip set. Despite being designed by different people, all of the chips are seen to be similar in size, area, and power dissipation. The cache controller and datapath chips are seen to be of comparable complexity (were the unnecessary columns removed from the tag RAM block this would be even more the case), while the cache RAM and instruction decoder, while being quite difficult in nature, are similar in size and complexity. This comparison suggests that it might be worthwhile in future designs to move some of the functionality of the cache controller into the instruction decoder.

⁴ The breakdown is only approximate; in many cases devices can be classified as belonging to several categories. Single and double emitter devices are counted as a single transistor. Diodes are not counted.

II.2.8. Communications

As the F-RISC/G prototype is partitioned, inter-chip communications becomes an important issue. Large fractions of the cycle time on are consumed by communication between chips.




Figure 77. Load critical path components.

Figure 78 shows a breakdown of the components of the LOAD critical path in the data cache, assuming that the Byte Operations chip is present. As can be seen, off-chip communications accounts for over 40% of the critical path. This is a unique design space that required special attention throughout the design process. Interestingly, these numbers are similar to those for the F-RISC/G adder critical path, as shown in Figure 79 adapted from [Phil93].




Figure 78. Components of adder critical path (adapted from [Phil93]).

Table II.2.8-1 lists the communications signals sent from the core CPU to the primary cache. Aside from an address and data, the CPU also sends out several handshaking and

control signals. These signals inform the caches of stalls and determine whether a LOAD or STORE is to take place.

Signal	Width	From	To	Description
ABUS	32	DP	DCC, ICC	Word (Instruction cache) or Byte (Data cache) address. Shared by both caches.
WDC	1	ID	DCC	Signals data cache to perform store.
STALLM	1	ID	DCC, ICC	Signals both caches to stall.
ACKI	1	ID	ICC	Signals instruction cache that it has caused a stall.
ACKD	1	ID	DCC	Signals data cache that it has caused a stall.
VDA	1	ID	DCC	Address on bus is valid for data cache.
IOCNTL	3	ID	DCC, ICC	Flush / Initialize / Write alignment
BRANCH	1	ID	ICC	Instruction cache should set RPC to address on bus.
DATAOUT	32	DP	DRAM	Word of data to be stored in data cache.

Table II.2.8-1. CPU to Cache Communications.

The IOCNTL lines are a 3-bit field that is part of the LOAD and STORE instructions, and are sent to both cache controllers. These bits are used to inform the caches when the system startup routine is complete, and to inform the data cache in the event of aligned byte or half-word writes. The meaning of the control bits are as shown in Table II.2.8.-1.

As the data cache receives a byte address from the datapath (unlike the instruction cache, which uses word addresses), support is provided using IOCNTL to allow reads and writes to any byte, half-byte, or word in the processor's address space. To read a non-word-aligned byte or half-byte, however, requires the presence of the Byte Operations chip on the MCM. Non-word aligned word-fraction Store support is provided in the DCC.

IOCNTL	Meaning
000	Read or write entire word
001	Read or write half-word
010	Read or write byte
011	Force a miss on this address
100	Co-processor support

Table II.2.8-2. IOCNTL Settings.

In order to prevent the need to design two different cache controllers, the cache controller chip is designed internally to handle word addresses. On the DCC, ABUS [2], the word address, must be wired to the pad ABUS [0]. Similarly, each bit on the bus is wired to the pad corresponding to its position in the word address. The two low order ABUS bits (byte address) are wired to the high order pads (See Figure 80). The controller chip knows to ignore these two bits when handling tags and presenting addresses to other chips, and uses them only when setting the RAMs into Write mode.




Figure 79. ABUS partitioning.

Table II.2.8-3 lists the signals sent from the cache to the CPU. These consist mostly of requested data, but also include signals to inform the CPU that a miss has occurred and the requested data will not be available in time.

Signal	Width	From	To	Description
MISSI	1	ICC	ID	A miss has taken place in the instruction cache.
MISSD	1	DCC	ID	A miss has taken place in the data cache.
INSTRUCTION	32	IRAM	ID	32-bit Instruction
DATAIN	32	DRAM	DP	Word of data for the datapath.

Table II.2.8-3. Cache to CPU communications.

II.2.9. Intra-cache Communications

The primary caches each consist of a single cache controller chip and eight cache RAM chips. While there is no inter-cache communication (i.e. the instruction and data caches do not communicate with each other), there is extensive communication between each cache controller and its associated RAM chips.

Table II.2.9-1 lists communications lines between the cache controllers and RAMs. The CRWRITE line is used to write into the cache RAMs. The CRWIDE line is used to toggle between the 4-bit per RAM CPU data path and the 64-bit per RAM L2 data path. The CRDRIVE line is used to control the bi-directional drivers / receivers used on the RAMs for communicating with the L2 cache.

Signal	MCM Length (mm)	Delay (ps)
ABUS	21	170
WDC	21	170
STALLM	26	190
ACKI	17	140
ACKD	25.5	190
VDA	21	170
BRANCH	15	135
DATAOUT	upper path: 22 lower path: 27	170 200
MISSI	18	150
MISSD	25	185
INSTRUCTION	fast bits: 13 slow bits: 24	120 170
DATAIN	upper path: 22 lower path: 28	170 200

Table II.2.9-1 MCM net lengths - CPU / cache signals.

Signal	Width	From	To	Description
CRABUS	9	CC	RAM	5-bit row address and 4-bit word address.
CRWRITE	4	DCC	DRAM	Write / Read .
HOLD	1	CC	RAM	Prevent RAM outputs from changing..
INLAT	1	CC	RAM	Allow 4-bit data input to pass through input latch.
CRWIDE	1	CC	RAM	Select wide input path (64-bit) for write from L2..
CRDRIVE	1	CC	RAM	Control bi-directional L2 bus.

Table II.2.9-2. Intracache communications.

Signal	Width	From	To	Description
L2ADDR	23	CC	L2	23-bit line address.
L2DONE	1	L2	CC	Indicates that the L2 has completed a transaction. Any data L2 places on the bus must be valid when this is asserted.
L2DIRTY	1	CC	L2	Indicates that the L2 will be receiving an address to be written into.
L2MISS	1	CC	L2	Indicates that the address on L2ADDR is needed by the CPU.
L2VALID	1	L2	CC	Indicates that the current data in the cache row specified by the cache tag currently being transacted is correct. De-asserted by L2 during TRAP.
L2SYNCH	1	CC	L2	A 1 GHz clock used for synchronizing with L2.
L2VDA	1	CC	L2	The address currently on L2ADDR is valid.

Table II.2.9-3. Secondary cache communications.

The HOLD and INLAT signals are used to latch the RAM 4-bit data outputs and inputs, respectively. The lengths of each of these lines or buses is less than 45 mm, for an estimated flight time of 300 ps.

II.2.10. Secondary Cache Communications

Table II.2.9-3 enumerates the signals used for communication between the primary and secondary caches.

Each cache controller will send out a 28-bit cache line address as soon as it is received from the CPU. This is done to allow the L2 cache to read its tag RAM simultaneously to the L1 cache. The cache controller will assert L2DIRTY as soon as it completes its tag RAM access if the accessed line is dirty. The L2 will not receive the address as stored in the primary caches tag RAM until later, however, and only if it is required (that is, a Stall occurs.)

The cache controller asserts L2MISS only if a miss occurs and the CPU acknowledges the miss. Whenever the address on the L2ADDR bus is valid, L2VDA is asserted.

Since the secondary caches do not have a synchronized clock, the L2SYNCH signal is used to inform the secondary caches that valid data is on the control and address lines. When the L2SYNCH signal goes high the data on the L2 communications lines is valid. It remains so for approximately 500 ps. If the MCM routing is done carefully, it may be possible to assure that the L2 communications signals are valid for as long as L2SYNCH is asserted.

The L2DONE signal is asserted by the L2 to indicate that it has performed the requested operations, both modifying its RAMs as appropriate and placing requested data on the bus. Any data being sent by the L2 must be on the bus for 750 ps prior to L2DONE being asserted.

In the event that the primary cache has to perform a copyback, the secondary cache will first receive the address (originating from the CPU and passing through the primary cache controller) that caused the copyback, along with the L2DIRTY signal and the data to be copied back, which should be latched at that point. Two more addresses will appear on the bus to the L2 (although they may or may not be valid), followed by address that had been stored in the tag RAM (the address of the data being copied back).

This "out of order" execution, in which the L2 may perform the read before the write on a copyback from the primary cache, allows maximum flexibility for the secondary cache designer (for example if two port RAM is available.)

II.2.11. Virtual Memory Support

The F-RISC/G CPU is designed with rudimentary support for virtual memory. Specifically, control and communications lines are provided to enable the caches to signal the CPU in the event of a page fault, as shown in Table II.2.11-1.

Signal	Width	From	To	Description
TRAPD	1	Cache	CPU	Data cache page fault
TRAPI	1	Cache	CPU	Instruction cache page fault
I1, I2, I3	3	Cache	CPU	Status lines sensed by PSW
O1, O2, O3	3	CPU	Cache	Status lines controlled by PSW

Table II.2.11-1. Virtual memory control.

The word addresses supplied by the CPU to the instruction cache and the byte addresses supplied by the CPU to the data cache are *virtual addresses* in that they refer to a location in the CPU's memory space without regard to their actual presence in physical memory. The CPU doesn't care where a particular virtual address maps to, as long as when data is requested from that address it is available.

Since the virtual instruction space is 2^{32} words in size and the data memory space is 2^{30} words in size, it is unlikely that the amount of physical RAM available in main memory will span the entire virtual memory space. In a typical virtual memory system, hardware and software is provided to allow the virtual memory to be divided into *pages* each of which may exist either in physical memory or on a secondary storage device, such as a disk drive. When the CPU requests a transaction to an address which is in a page not currently in physical memory, a *page fault* occurs, and the page which is needed is loaded from secondary storage, replacing another page already in physical memory if necessary. Since the amount of time necessary to access the secondary storage device, transfer the existing memory page to this device, locate the required page on the disk, and retrieve it back into memory is extremely long compared to the CPU cycle time⁵, it is desirable for the cache to inform the CPU of the problem and allow the CPU to proceed with other instructions while the page swap occurs, if possible. This is typically performed by the operating system which will *context switch* to another waiting, unrelated process.

Due to the hardware cost of such a system, the virtual to physical address translation cannot occur in the primary cache. Instead, it is expected that some higher level of memory, perhaps the level just before main memory, will handle the translation of virtual addresses into physical addresses. When a page fault occurs at this level of memory, the

⁵ The access time for a typical hard disk drive is on the order of 10 ms, or one million CPU cycles.

CPU is informed via the TRAPD or TRAPI signal. The CPU then handles the interrupt by branching to the appropriate trap vector. It is presumed that the operating system has installed code at the appropriate trap vector to handle page faults. The caches will send "DONE" signals all the way down to the primary cache, which will recover from its stall and lower the MISS line as if it had the correct data. The cache must then be re-validated through a flush of the incorrect address. The CPU will lower the STALL and ACK in response to the primary cache lowering its MISS, and will prevent it from going high again in response to the incoming TRAP.

Typically, the CPU, upon receiving the TRAP, will perform instructions which don't involve the memory location which page faulted, and, when the page is finally available, will re-issue the request. The CPU contains pipeline stages, which enable it to re-issue a LOAD or STORE which result in a page fault.

The exact behavior of the CPU in response to a memory page fault depends on the contents of the CPU pipeline and the state of the caches at the time the page fault occurs.

	Abort DW?
Reset	YES
System error	YES
Data cache page fault	YES
Arithmetic trap	NO
Software trap	NO
Instruction cache page fault	NO
Device interrupt	NO
User interrupt	NO

Table II.2.11-2. CPU trap behavior.

II.2.12. Timing

As mentioned earlier, the cache memory hierarchy has its own critical paths. The most critical of these is the path from address generation at the CPU to data reception by the CPU.

II.2.12.1. Load Timing

Figure 81 is a timing diagram of data cache LOAD operations. This timing diagram is based on the back-annotated (post-route) netlists for the cache controller, instruction decoder, and datapath chips. The vertical timing lines represent synchronized clock phase 1. Slightly after phase 1 of the first cycle, the CPU puts address (20)_{hex} on the ABUS (Table II.2.12.1-1).

It arrives at the data cache controller during phase 2 where it passes through the master of pipeline latch 1. The WDC and VDA lines are stable prior to the address. On the DCC, the tag RAM receives its inputs (address and data) from the master of pipeline latch 1, while the slave is used to feed the comparator. The tag RAM read access time is approximately 500 ps.

Signal	Delay from phase 1
ABUS	145
WDC	75
VDA	-100
BRANCH	-65
DOUT	210


Table II.2.12.1-1. Back-annotated signal timings.

After the cache RAMs supply the data to the CPU, the only remaining task for the cache is to inform the CPU that the data is available and to re-synchronize with the CPU's pipeline.

The situation is more complicated if the cache row corresponding to the cache access is marked as dirty. If a miss occurs and the cache row is dirty, the primary cache must send the current contents of that row to the secondary cache before overwriting it with the data requested by the CPU.

Figure 82 is an example of code that would result in this condition. The first line of code stores the contents of register 2 into cache row 2 (the row is calculated by bits 4 through 8 of the address). The corresponding tag would be 0, and the dirty bit would be set to indicate that the CPU has changed the contents of this address and that the higher levels of memory are out of date.

The two ADDI instructions are used to set register 3 to 3FFFFFFE20 (the use of two instructions is necessary since no F-RISC instructions accept 32-bit literal values). Finally, the LOAD instruction should fetch the contents of 3FFFFFFE20 into register 1.



```

ADDI R3=0+FE20      ; the add instructions are used to
ADDI R3=0+3FFF /LDH ; assemble 3FFFFE20 as the destination for the
LOAD                ;
LOAD R1=[0+R3]      ; put the contents of 1024 into R1

```

Sample LOAD copyback code fragment

3FFFFE20 corresponds to cache row 2 and tag value 1FFFFF. Since row 2 previously held tag 0, a miss will occur. Since the dirty bit for row 2 is set, a copyback must first take place.

Figure 82 is the timing diagram for this example. The STORE request is received by the primary cache at time 9375. In order to show the worst case, only one cycle of latency is allowed on this timing diagram between the STORE and subsequent LOAD. The LOAD request is received at time 11375.

II.2.13. Store Timing

Figure 83 is a timing diagram showing consecutive STORE instructions. When a STORE is to take place, the instruction decoder signals the cache controller by asserting the WDC signal. Since the signal is derived from the instruction word and can be sent directly from the instruction decoder rather than the datapath chips, the signal arrives a few hundred picoseconds before the address (at time 9075 in this example).

Every STORE instruction is allocated two cycles by the CPU. The second cycle is necessary because a STORE requires a read from and a write to the tag RAM.

For the first of the two cycles, the cache controller will be in the READ state. While in this state, the cache controller checks the tag RAM in order to determine whether a hit has occurred. As far as the cache controller is concerned, the first half of a STORE instruction proceeds identically to a LOAD instruction.

The cache controller latches the address from the CPU during the first half of the STORE, so the CPU does not have to keep the address stable for two cycles. During the second cycle the comparator calculates the result.

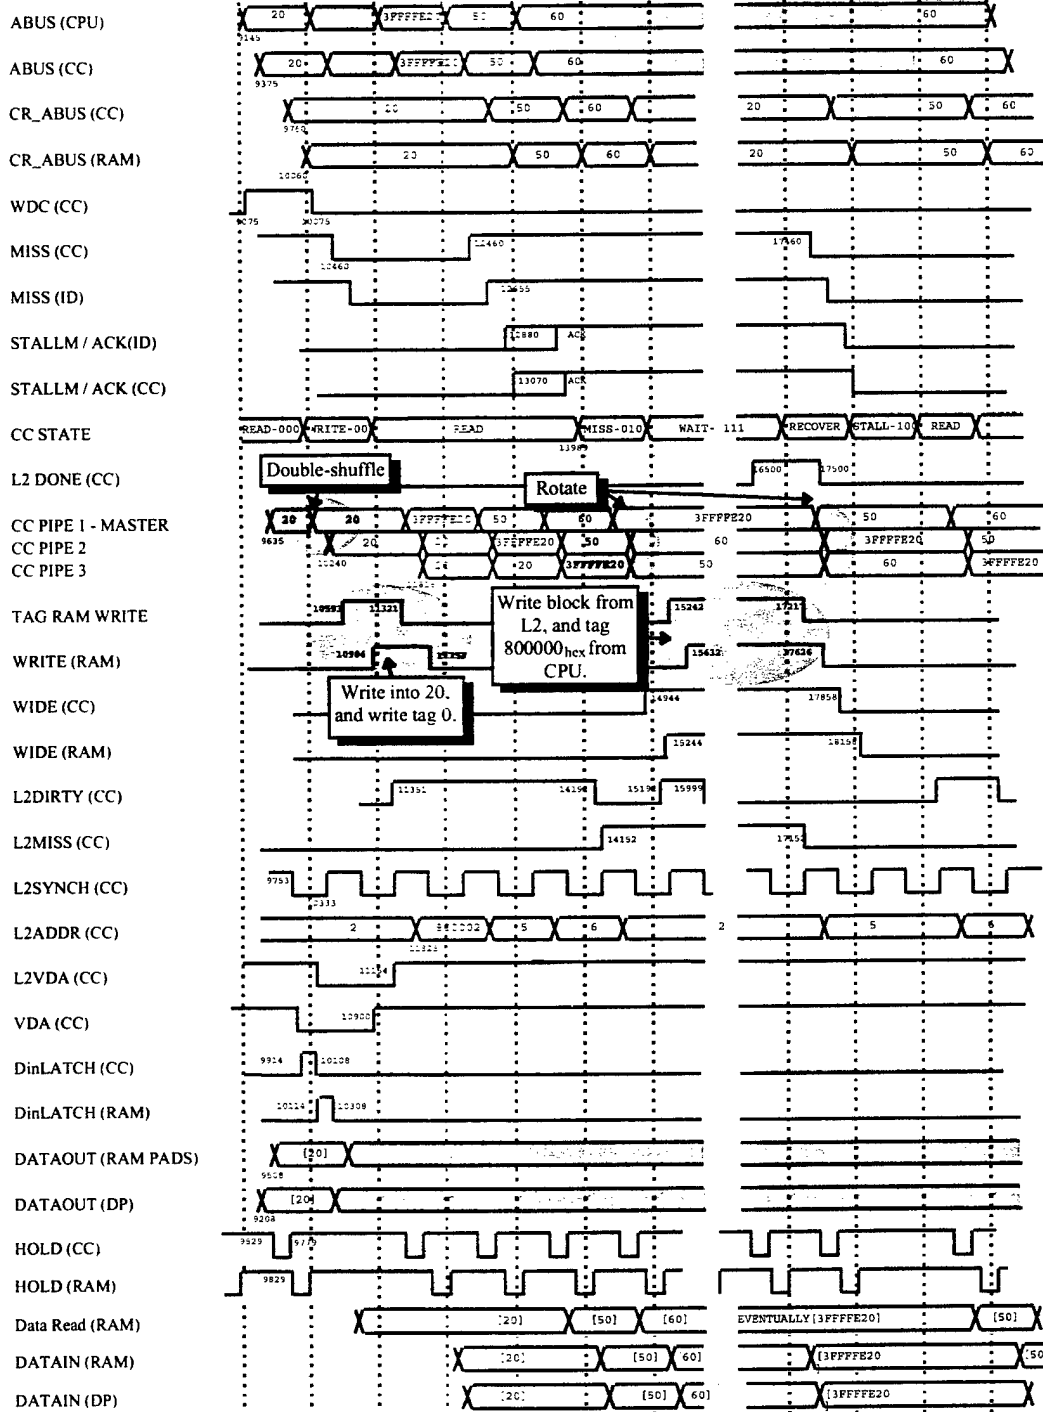


Figure 81. Data Cache Timing - Load Copyback.




Figure 82. Instruction Cache Miss Timing.

II.2.14. Instruction Fetch Timing

The instruction cache timing is, in most respects, similar to the timing of the data cache during a LOAD. This is particularly true when a BRANCH occurs.

The instruction cache controller contains a remote program counter (RPC) that is used to generate addresses to fetch and send to the CPU. This occurs without any intervention from the datapath or instruction decoder. In the event of a BRANCH, the address is received off of the ABUS, as in the data cache.

Unlike in the data cache, it is not necessary to delay the data sent to the CPU using the HOLD signal, since the instruction cache timing is much more constrained.

When the CPU starts up, a "phantom" BRANCH to location 20_{hex} is injected into the pipeline. Figure 84 illustrates how such a BRANCH might take place. As in the data cache, the target address is expected to be available at the cache controller at approximately 375 ps after "phase 1" (simulation time 9375). The actual BRANCH signal arrives approximately a phase earlier.

The timing of the instruction cache is more critical than that in the data cache. The architecture was designed to support a byte-operations chip in the data cache; by not including it, the timing in the data cache became fairly relaxed. The instruction cache has only from 1850 ps - 2100 ps in which to perform a fetch, versus 2250 ps in the data cache. Bits 3-7 of the instruction word must arrive at the instruction decoder a phase earlier than the remaining 27 bits.

In order to allow bits 3-7 (the "fast" bits) to arrive more quickly, the two RAMs, which supply these bits to the instruction decoder, were placed as close to the ID as possible without increasing the distance from the instruction cache controller.

If the CPU determines that the request to the cache cannot be flushed, it must stall, and will assert the STALLM line, which is shared by both caches.

Upon receiving STALLM, each cache will move into the MISS state. At the time this occurs, neither cache knows whether or not it is the cache, which caused the stall. In order to inform the appropriate cache that it is responsible for the stall, the instruction decoder will assert the appropriate acknowledgment line (either ACKI or ACKD).

The cache that receives both the ACK and the STALLM will progress through the normal miss cycle as previously described. The other cache will behave almost identically, but

will skip the WAIT state, thus preventing any cache state information from being overwritten. This cache will skip directly into the RECOVER state, and, once cycle later, will enter the STALL state where it will idle while awaiting STALLM to be de-asserted. Since the pipeline rotate occurs only in the RECOVER state (rather than in the STALL state), the pipeline in the non-stalled cache will be correct for when the CPU recovers from the stall.

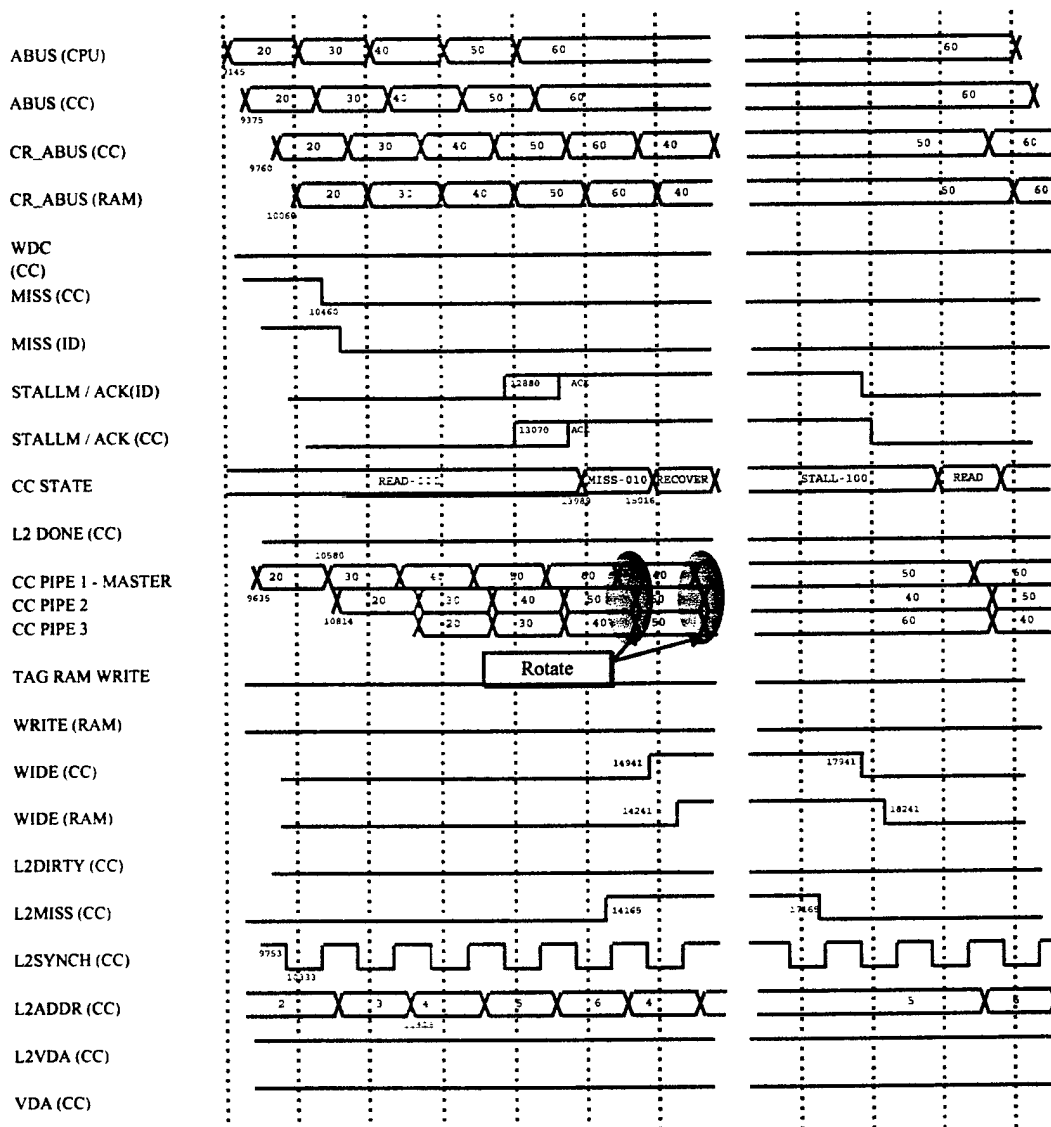


Figure 83. Data Cache During Instruction Cache Stall.

II.2.15. Other Cache Stalled

When a cache determines that a miss has occurred and that it will not be able to satisfy the CPU's request in the time allotted, the cache controller will assert the appropriate MISS line (MISSI for the instruction cache, or MISSD for the data cache).

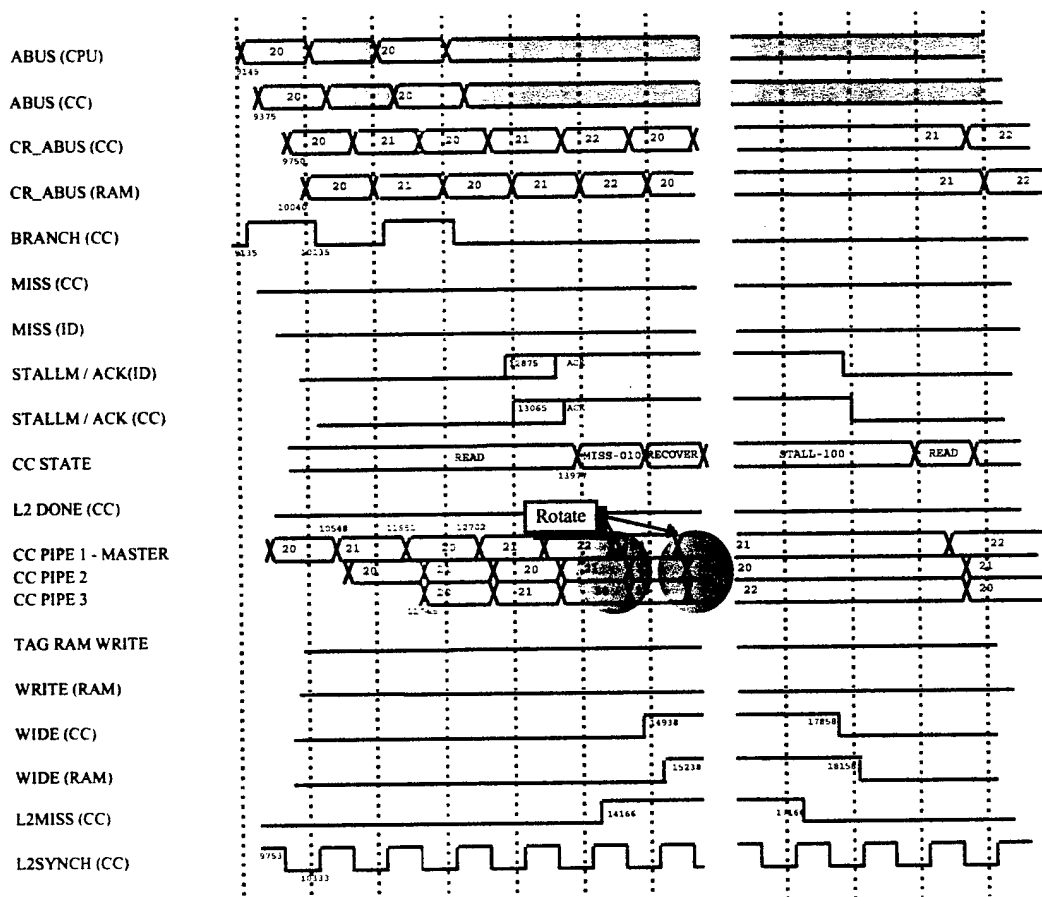


Figure 84. Instruction cache during a data cache stall.

II.2.16. Processor Start-up

One of the most important responsibilities of the cache is to enable the processor to correctly start up. When the processor is powered on, or reset, it needs to be fed the appropriate start-up instructions, and the data cache must be invalidated or pre-loaded with valid data.

When the processor is initialized, it inserts an unconditional BRANCH to location 20_{hex} into the pipeline. It is the responsibility of the instruction cache to fetch this instruction upon receiving the BRANCH signal and the address.

Figure 85 illustrates the timing at the instruction cache controller during processor start-up. The cache controller will receive the branch request and must realize that a miss must occur, regardless of whether the tag in the tag RAM accidentally matches the tag of the start-up address (0). This is accomplished through coordination with the secondary cache, since too little handshaking exists between the CPU and the cache to enable this to be self-contained.

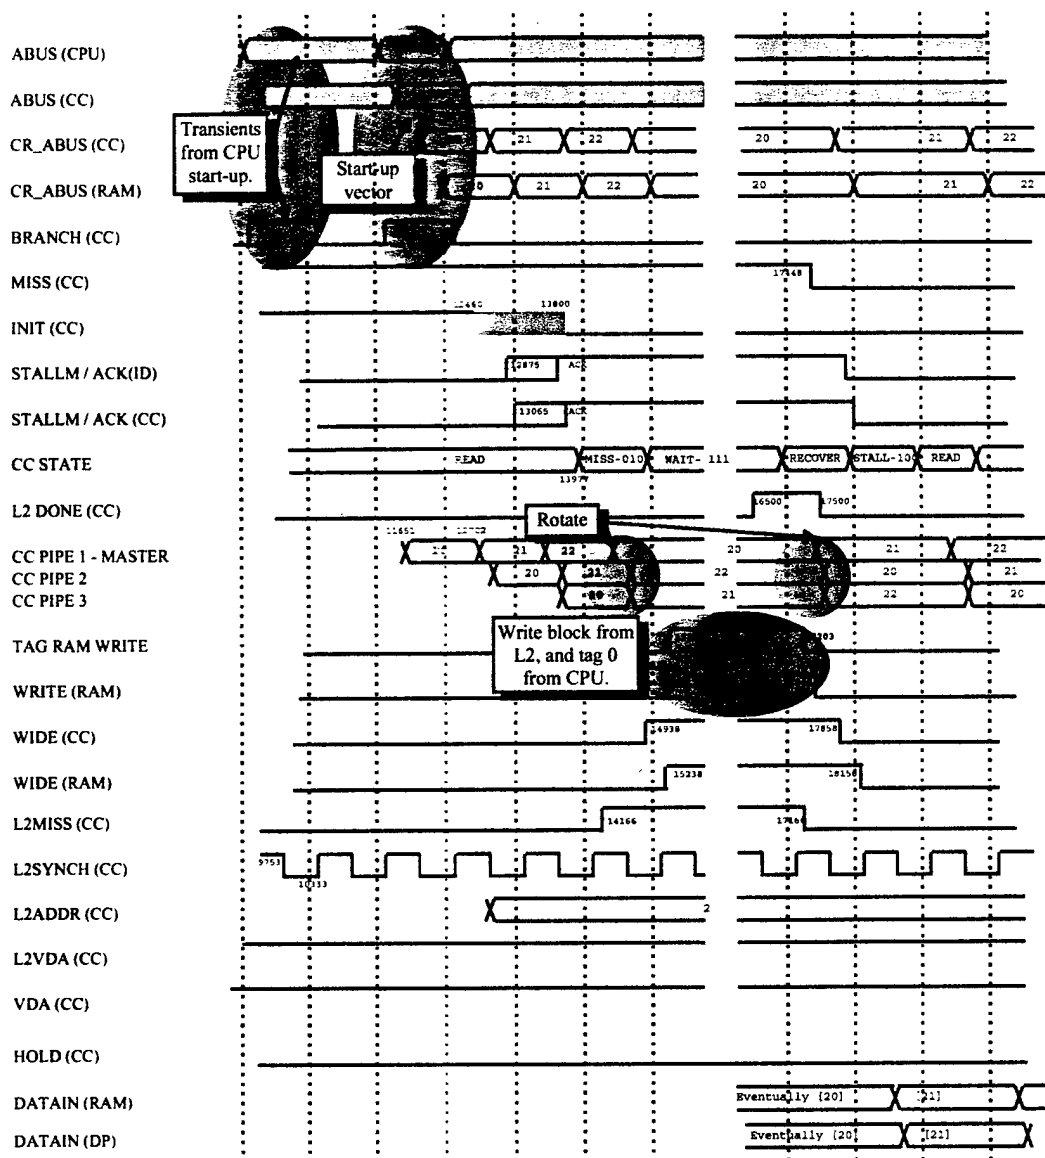


Figure 85. Instruction cache at start-up.

The secondary cache will receive the global RESET line (as well as all external trap and interrupt lines) and is responsible for initializing the CPU and the cache in the proper sequence.

Figure 86 illustrates the operation of the instruction cache during a page fault or during a trap, which happens to occur coincidentally to a secondary cache transaction. The cache must take special measures to preserve the integrity of the tag RAM during such an event. When a page fault occurs, at least once primary cache (the one corresponding to the fault) is awaiting data from the secondary cache.

The primary cache will be in the WAIT state, with the tag RAM and cache RAM WRITE signals asserted. The cache RAMs will be performing a wide WRITE, awaiting the data from the secondary cache. The tag RAM will be writing in the new tag from the pipeline (originating from the CPU) along with the appropriate value of DIRTY. The old tag will have already been sent to the secondary cache during the READ stage of that memory access cycle.

When the trap occurs (presumably at the main memory level of the memory hierarchy in the case of a page fault), the trap is sent to the secondary cache. The secondary cache will then de-assert the L2VALID line. This bit is stored in the appropriate row of the tag RAM, along with the appropriate tag. If the bit is set to "valid," then future cache operations on that tag will proceed as normal. If, however, the data transfer from the secondary cache is interrupted by a trap, then the secondary cache sets the bit to "invalid," and if another operation takes place on that tag, it automatically causes a miss to take place.

In the event that a STORE into the data cache caused the page fault, it is questionable as to whether the transaction should be interrupted. If the cache were to simply mark the tag as invalid, the data stored by the CPU would be lost, and the CPU would have no way of knowing about it. Since STOREs are comparatively rare, and STORE misses even more so, the best decision is simply to stall the processor until the primary cache has valid data.

Since it takes approximately 750 ps to write into the tag RAM, and the data should be stable for a considerable period before that, the secondary cache should wait two cycles after deasserting L2VALID before sending the trap signal through to the primary caches and CPU.

The primary cache responds to the trap signal by resetting to the READ state. The MISS signals may be spuriously asserted by the primary cache while the trap is held high (the trap is tied to the INIT signal pad), but the secondary cache has enough information to ignore it, and the CPU ignores misses which occur while processing the trap.

Figure 87 illustrates the interaction of the F-RISC / G caches during a load copyback. The primary cache sends an address to the secondary cache before it is determined whether the primary cache needs the address. By the time the miss signal is sent to the secondary cache, assuming the secondary cache has not received additional valid addresses (the primary cache will assert the L2VDA signal when a valid address is on the bus), the secondary cache has already had at least a cycle to perform a read. The secondary cache must finish the read, and, using the copyback address and data, which is sent to the secondary cache following the L2MISS signal, perform a write. While the write is being performed, the data read from the secondary cache must be latched. Once the data on the bidirectional bus is no longer needed, the secondary cache can assert

L2DONE signal and put the data on the bus (the data should be on the bus for a phase before L2DONE is asserted.)




Figure 87. Load copyback in F-RISC / G cache.

It is important to note that the five cycle mean access time for the secondary cache was based on calculations for the stall component of CPI. Therefore, the required five-cycle limit implies that, on average, accesses to the secondary cache result in a stall of only five cycles. Since, in the event of a primary cache hit, the data is required at the CPU at approximately the same time the secondary cache receives the miss signal in the event of a primary cache miss, the five cycles allotted to the secondary cache begin approximately when the secondary cache receives the L2MISS signal. This means that, on average, a primary cache read miss has 7 ns to be completed. (The data cache has an additional phase, while the instruction cache fast bits have one phase fewer).

II. LIST OF ALL PUBLICATIONS AND TECHNICAL REPORTS

- [1] **"Cell Library for Current Mode Logic using an Advanced Bipolar Process,"** (J. F. McDonald, H. J. Greub, T. Yamaguchi, and T. Creedon), I.E.E.E. J. Sol. State Cir.. Special issue on VLSI, (D. Bouldin, guest editor), I.E.E.E. Trans. on Solid State Circuits, Vol. JSSC-26(#5), pp. 749-762, May, 1991.

- [2] **"F-RISC/I: Fast Reduced Instruction Set Computer with GaAs H-MESFET Implementation,"** Proc. I.E.E.E. Int. Conf. on Computer Des., (J. F. McDonald, C. K. Tien, C. C. Poon, H. Greub) Boston, MA, (I.E.E.E. Cat. # CH3040-3/91/0000/0293), pp. 293-296, October 14-16, 1991.

- [3] **"F-RISC/G: AlGaAs/GaAs HBT Standard Cell Library,"** Proc. I.E.E.E. Int. Conf. on Computer Des., (J. F. McDonald, K. Nah, R. Philhower, J. S. Vanetten, S. Simmons, V. Tsinker, Maj. J. Loy, and H. Greub), Boston, MA, (I.E.E.E. Cat. # -3/91/0297), pp. 297-300, October, 1991.

- [4] **"Wideband Wafer-Scale Interconnection in a Wafer Scale Hybrid Package for a 1000 MIPS Highly Pipelined GaAs/AlGaAs HBT Reduced Instruction Set Computer,"** Proc. 1992 Int. Conf. on Wafer Scale Integration, ICWSI-4, San Francisco, January 20, 1992, Reprinted Hardbound by Computer Science Press, V. K. Jain, and P. W. Wyatt, Eds. [I.E.E.E. CS#2482], pp. 145-154. (J. F. McDonald, R. Philhower, J. S. Van Etten, S. Dabral, K. Nah. and H. Greub).

- [5] **"Bypass Capacitance for WSI/WSHP Applications,"** Proc. Fifth Int. Conf. on WSI, ICWSI93, San Francisco, CA, M. Lea, Ed., I.E.E.E. Computer Soc. Press, pp. 218-228, February, 1993 (J. F. McDonald, H. Greub, R. Philhower, J. Van Etten, K. S. Nah, P. Campbell, C. Maier, Lt. C. J. Loy, P. Li, L. You, and T.-M. Lu).

- [6] **"Fluorinated Parylene as an Interlayer Dielectric for Thin Film MultiChip Modules,"** spring 1992 meeting of the Materials Research Society, Reprinted in Vol. 264 of the MRS Symposium Proceedings, Electronic and Packaging Materials Science VI, Paul S. Ho, K. A. Jackson, C.-Y. Li and G. F. Lipscomb, Eds., pp. 83-90, 1993 (J. F. McDonald, S. Dabral, X. Zhang, W. M. Wu, G.-R. Yang, C. Lang, H. Bakhru, R. Olsen and T.-M. Lu)

- [7] **"A 500ps 32 X 32 Register File Implemented in GaAs/AlGaAs HBT's,"** Proc. I.E.E.E. GaAs Symposium [I.E.E.E. Cat. 93CH3346-4], San Jose, Oct. 1993, pp. 71-75, (J. F. McDonald, K. S. Nah, R. Philhower, and H. Greub).

- [8] **"F-RISC/I: A 32 Bit RISC Processor Implemented in GaAs H-MESFET Super Buffer Logic,"** Proc. I.E.E.E. GaAs Symposium [I.E.E.E. Cat. #93CH3346-4], San Jose, CA, Oct. 1993, pp. 145-148, (J. F. McDonald, C. K. Tien, K. Lewis, R. Philhower, and H. J. Greub).

- [9] **"Frequency Domain (1kHz-40GHz) Characterization of Thin Films for Multichip Module Packaging Technology,"** (J. F. McDonald, W.-T. Liu, S. Cochrane, X.-M. Wu, P. K. Singh, X. Zhang, D. B. Knorr, E. J. Rymaszewski, J. M. Borrego, and T.-M. Lu), Elect. Lett., Jan. 20, 1994, Vol. 30(#2), pp. 117-118.

- [10] **"Poly-tetrafluoro-p-xylylene as a Dielectric for Chip and MCM Applications,"** (J. F. McDonald, S. Dabral, G.-Y. Yang, X. Zhang, & T.-M. Lu, J. Vac. Sci. and Technol., B 11(#5), Sept./Oct. 1993, pp. 1825-1830.

- [11] **"Application of a Floating-Random-Walk Algorithm for Extracting Capacitances in a Realistic HBT Fast-RISC RAM Cell,"** (J. F. McDonald, Y. L. Le Coz, R. B. Iverson, H. J. Greub, P. M. Campbell, and J. F. McDonald), Proc. I.E.E.E. VLSI Multi-Layer Interconnect Conf., V-MIC94, Santa Clara, CA. June, 1994, pp. 542-544.

- [12] **"Design of a Package for a High Speed Processor Made with Yield Limited Technology,"** (J. F. McDonald, A. Garg, J. Loy, and H. Greub), Proc. I.E.E.E. Fourth Great Lakes Symposium on VLSI, March 4-5, 1994, Notre Dame University, Indiana, [I.E.E.E. Cat. #94TH0603-1, Comp. Soc. # 5610-02], pp. 110-113.

- [13] **"Wiring Pitch Integrates MCM Wiring Domains,"** (J. F. McDonald, J. Loy, A. Garg, M. Krishnamoorthy), Proc. I.E.E.E. Fourth Great Lakes Symposium on VLSI, March 4-5, 1994, Notre Dame University, Indiana, [I.E.E.E. Cat. #94TH0603-1, Comp. Soc. # 5610-02], pp. 110-113.

- [14] **"Differential Routing of MCMs — CIF: The Ideal Bifurcation Medium,"** (J. F. McDonald, J. Loy, A. Garg, M. Krishnamoorthy), Proc. I.E.E.E. Int. Conf. on Computer Des., Cambridge, MA, [I.E.E.E. Cat. # 94CH35712], pp. 599-603, October 10-12, 1994.

- [15] **"Thermal Design of an Advanced Multichip Module for a RISC Processor,"** (J. F. McDonald, A. Garg, J. Loy, H. Greub, T.-L. Sham), Proc. I.E.E.E. Int. Conf. on Computer Des., Cambridge, MA, [I.E.E.E. Cat. # 94CH35712], pp. 608-611, October 10-12, 1994.

- [16] **"Three Dimensional Stacking with Diamond Sheet Heat Extraction for Subnanosecond Machine Design,"** (J. F. McDonald, H. Greub, A. Garg, P. Campbell, S. Carlough, C. Maier), Proc. 1995 Int. Conf. on Wafer Scale Integration, ICWSI-7, San Francisco, January 20-22, 1995, Reprinted in Hardbound by Society Press, S. K. Tewksbury, and G. Chapman, Eds. [I.E.E.E. CS #2482], pp. 62-71.

- [17] **"Design of a 32-bit Monolithic Microprocessor Based on GaAs H-MESFET Technology,"** in review for I.E.E.E. Transactions on VLSI Systems," (J. F. McDonald, C.-K. V. Tien, K. Lewis, H. J. Greub, and T. Tsen).

- [18] **"A Very Wide Bandwidth Digital VCO Implemented in GaAs HBT's Using Frequency Multiplication and Division,"** (J. F. McDonald, P. M. Campbell, H. J. Greub, A. Garg, S. Steidl, C. Maier, and S. Carlough), Proc. 17th Ann. I.E.E.E. GaAs Symposium, San Diego, CA, October 29-November 1, 1995.

- [19] **"Metal-Parylene Interaction Systems,"** (J. F. McDonald, S. Dabral, X. Zhang, B. Wang, G.-R. Yang, and T.-M. Lu), Mat. Res. Soc. Proc., Vol. 381, in "Low K Dielectrics," S. Murarka, T.-M. Lu, T.-S. Kuan, and H. C. Ting, Eds., pp. 205-218, 1995.

- [20] **"Low Dielectric Constant Polymers for On-Chip Interlevel Dielectrics with Copper Metallization,"** (J. F. McDonald, R. J. Gutmann, T. Paul Chow, D. J. Duquette, T.-M. Lu, and S. P. Murarka), Mat. Res. Soc. Proc., Vol. 381, in "Low K Dielectrics," S. Murarka, T.-M. Lu, T.-S. Kuan, and H. C. Ting, Eds., pp. 177-195.

- [21] **"Unterminated Bonds in Parylene-N Films,"** (J. F. McDonald, X. Zhang, B. Wang, and S. Dabral), Semiconductor International, Vol. 14, December 1995, pp. 89-94.

- [22] **"Crystallinity Properties of Parylene N Affecting ILD Applications,"** Paper H2.11 at the ICMCTF 95, reprinted in Thin Sol. Films, Vol. 270, (1-2), pp. 508-511, December 1, 1995.

- [23] **"Improvement of Parylene N Deposition Rate by Electric Field,"** I.E.E.E. Dielectrics for Ultralarge Scale Integration Multilayer Interconnection Conference, DUMIC 96, Santa Clara Marriott Hotel, Santa Clara, CA, pp. 214-221. February 20-21, 1996 (J. F. McDonald, B. Wang, T.-M. Lu, and G. Yang)

- [24] **"Very Fast RISC Processors - Subnanosecond Computing using High Performance HBT Devices and MCM packaging,"** Gov. Microcir. Applic. Conf., GOMAC96, Orlando, Hyatt Orlando Inn, Orlando, Florida, March 29-31, 1996, pp. 217-221, (J. F. McDonald and H. Greub).

- [25] **"Chip Pad Migration is a Key Component in High Performance MCM Design,"** (J. F. McDonald, J. Loy, A. Garg, and M. Krishnamoorthy), Sixth Great Lakes Symposium on VLSI, GLSVLSI96, (I.E.E.E. Cat. 96TB100041), Iowa State University, March 22-23, 1996, pp. 96-99.

- [26] **"Fast RISC Design using 100 GHz HBT and Microwave HDI MCM Technology,"** (J. F. McDonald and H. Greub), RADC Workshop on Academic Electronics in New York State, Embassy Suites Hotel, Syracuse, New York, June 13-14, 1996, pp. 259-264.

- [27] **"Dual Damascene Structure Fabrication with Parylene-N as the ILD and Copper as the Interlayer Metal,"** (J. F. McDonald, B. Wang, C. Steinbruchel, and R. Tacito), 13th annu. VLSI MultiLevel Interconnection Conference, VMIC96, Santa Clara Marriott, Santa Clara, CA, June 18-20, 1996, pp. 58-60.

- [28] **"A Floating Random Walk Method for Computing Interconnection Capacitances,"** (J. F. McDonald, Y. L. LeCoz, H. J. Greub, A. Garg, and R. Iverson), 13th annu. VLSI MultiLevel Interconnection Conference, VMIC96, Santa Clara Marriott, Santa Clara, CA, June 18-20, 1996, pp. 230-232.

- [29] **"Accurate High-Speed Performance Prediction for Full Differential Current-Mode Logic: The Effect of Dielectric Anisotropy,"** (J. F. McDonald, A. Garg, Y. L. LeCoz, H. J. Greub, R. B. Iverson, R. F. Philhower, P. M. Campbell, C. A. Maier, S. A. Steidl, M. W. Ernest, R. P. Kraft, S. R. Carlough, J. W. Perry, T. W. Krawczyk, I.E.E.E. Transactions on Computer Aided Design of Integrated Circuits and Systems, Vol. 18 (#2), February, 1999, 212-219.

- [30] **"Optimal Differential Routing based on Finite State Machine Theory,"** (with M. S. Krishnamoorthy, and James Loy), J. VLSI Design, Vol. 9, 1999, (# 2), pp. 105-117.

- [31] **"Study of Fluorine Diffusion in Metallized Polymers using Ion Beam Techniques,"** (with A. Kumar, H. Bakhru, B. Wang, G.-R. Yang, J. Fortin, and T.-M. Lu), *J. Materials, Chemistry and Physics*, Vol. 59, June 1999, pp. 136-138.

**III. LIST OF ALL SCIENTIFIC PERSONNEL SHOWING
ADVANCED DEGREES EARNED BY THEM WHILE
EMPLOYED ON THE PROJECT FROM ITS INCEPTION
(1991)**

- [1] Lt. Cmdr. James Loy, "Differential Routing Tools for High Speed GaAs HBT CML Circuits," Ph.D. 1993.

- [2] Robert Philhower, "Spartan RISC Architecture for Yield Limited Technology," Ph.D. 1993.

- [3] Kyung Suc Nah, "An Adaptive Clock Deskew Scheme and a 500 ps 32 by 8 Bit Register File for a High Speed Digital System," Ph.D. 1994.

- [4] C.-K. Vincent Tien, "System Design, Analysis, Implementation and Performance Evaluation of a 32 Bit RISC Processor Based on GaAs HMESFET Technology," Ph.D. 1994.

- [5] Cliff Maier, "High Speed Microprocessor Cache Memory Hierarchy for Yield Limited Technology," Ph.D. August, 1996.

- [6] Atul Garg, "Characterization of AlGaAs/GaAs HBT Interconnections for Digital Applications, and Design of a Dense MCM for a (2GHz Clock) 1 GOPS Fast RISC Processor," Ph.D. August 1997.

- [7] Peter Campbell, "Optimization of Circuits and Physical Design in a High Speed Bipolar RISC Processor," Ph.D. August 1997.

IV. LIST OF INVENTIONS BY NAME

No formal patent applications have been filed during this grant due to lack of funds for legal expenses

V. REFERENCES

- [1]. R. C. Eden, "Comparison of GaAs Device Approaches for Ultrahigh-Speed," Proc. of the I.E.E.E., Vol. 70,(#1), January 1982, pp. 5-12.
- [2]. H. Kroemer, "Heterostructure Bipolar Transistors and Integrated Circuits, Proc. of the I.E.E.E., Vol. 70,(#1), January 1982, pp. 13-26.
- [3]. H. Kroemer, "Theory of Wide-Band Gap Emitter for Transistors," Proc. I.R.E., Vol. 45(#11), pp. 1135-1137. Nov. 1957.
- [4]. Claim 2 of US Patent 2,569,347 to W. Shockley, Filed 26, June 1948, Issued, 25, Sept. 1951, expired Sept. 1968.
- [5]. W. P. Dumpke, J. M. Woodall, and V. L. Rideout, "GaAs-AlGaAs Heterojunction Transistor for High Frequency Operation," Solid State Electronics, Vol. 15(#12), pp. 1334-1349, December 1972.
- [6]. M. F. Chang, Ed., "Current Trends in Heterojunction Bipolar Transistors," Selected Topics in Electronics and System - Vol. 2, World Scientific Press, ISBN # 981-02-2097-9, 1996.
- [7]. X. Tang, V. K. De, and J. D. Meindl, "Intrinsic MOSFET Parameter Fluctuations Due to Random Dopant Placement," I.E.E.E. Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 5(#4), December 1997, pp. 369-376.
- [8]. K. Suzuki, and T. Sugii, "Analytical Models for n⁺p⁺ Double-Gate SOI MOSFET's, I.E.E.E. Transactions on Electron Devices, Vol. 42(#11), November, 1995, pp. 1940-1947.
- [9]. T. Tanaka, H. Horie, and T. Sugii, "Ultrafast Operation of V_{th}-Adjusted p⁺-n⁺ Double-Gate SOI MOSFET's," "Ultrafast Operation of V_{th}-Adjusted p⁺-n⁺ Double-Gate SOI MOSFETS," I.E.E.E. Electron Device Letters, Vol. 15(#10), October 1994, pp. 386-387.
- [10]. B. S. Meyerson, "UHV-CVD SiGe Technology," Proc. I.E.E. E. Vol. 80(#9), Sept. 1992, pp. 1592-1608.
- [11]. J. A. Simmons, M.A. Blount, J. S. Moon, W. E. Baca, J. L. Reno, and M. J. Hafich, "Unipolar Complementary BiD Quantum Tunneling Transistor, I.E.E.E. Electronic Devices Meeting, Washington, DC, Dec. 7-10, pp. 755-758.
- [12]. Y. Taur, and T.-H. Ning, "Fundamentals of Modern VLSI Devices," Cambridge University Press, ISBN 0-521-55056-4, NY, 1998.

- [13]. M. Yang, C.-L. Chang, M. Carroll, and J. C. Sturm, "25nm p-Channel Vertical MOSFET's with SiGeC Source-Drains," I.E.E.E. Electron Device Letters, Vol. 20(#6), June 1999, pp. 301-303.
- [14]. L. D. Lanzerotti, A. St. Amour, C. W. Liu, J. C. Sturm, J. K. Watanabe, and N. D. Theodore, "Si/Si(1-x-y)Ge(x)C(y)/Si Heterojunction Bipolar Transistors," I.E.E.E. Electron Device Letters, Vol. 17(#7), July 1998, 334-337.
- [15]. D. J. Roulston, "Bipolar Semiconductor Devices," McGraw-Hill Publishing, ISBN 0-07-054120-5, NY, 1990.
- [16]. P. M. Solomon and D. D. Tang, "Bipolar Circuit Scaling," Proc. 1979 International Solid State Circuits Conference, ISSCC Digest of Technical Papers, pp. 86-87, 1979.
- [17]. D. L. Harame, J. H. Comfort, J. D. Cressler, E. F. Crabbe, J. Y.-C. Sun, B. S. Meyerson and T. Tice, "Si/SiGe Base Epitaxial Transistors, Part I, Materials, Physics and Circuits," I.E.E.E. Transactions on Electronic Devices, Vol. 42, pp. 452-468, 1995.
- [18]. D. L. Harame, J. H. Comfort, J. D. Cressler, E. F. Crabbe, J. Y.-C. Sun, B. S. Meyerson and T. Tice, "Si/SiGe Base Epitaxial Transistors, Part II, Process Integration and Analog Applications," I.E.E.E. Transactions on Electronic Devices, Vol. 42, pp. 469-482.
- [19]. K. Washio, E. Ohue, H. Shimamoto, K. Oda, R. Hayami, Y. Kiyota, M. Tanabe, M. Kondo, T. Hashimoto, and T. Harada, "A 0.2 micron 180 GHZ fmax, 6.7-ps-ECL SOI/HRS Self-Aligned SEG SiGe HBT/CMOS Technology for Microwave and High-Speed Digital Applications, IEDM, 2000, pp. 741-744.
- [20]. C. Y. Chang, and Francis Kai, GaAs High-Speed Devices, John Wiley, 1994.
- [21]. R. Anholt, Electrical and Thermal Characterization of MESFETs, HEMTs, and HBTs, Artech House, 1995.
- [22]. B. Jalali, and S. J. Pearson, Eds., InP HBTs, Growth, Processing and Applications, Artech House, 1995
- [23]. R. Williams, Modern GaAs Processing Techniques, Artech House, 1991
- [24]. U. Ciligioglu, Systematic Analysis of Bipolar and MOS Transistors, Artech House, 1994.
- [25]. F. Ali, and A. Gupta, Eds., HEMTs & HBTs, Artech House, 1991.
- [26]. J. W. Mayer and S. S. Lau, Electronic Materials Science for Integrated Circuits in Si and GaAs, Macmillan, 1990.

- [27] N. Kanopoulos, Gallium Arsenide Digital Integrated Circuits, Prentice Hall, 1989.
- [28] S. Long, and S. Butner, Gallium Arsenide Digital Integrated Circuit Design, McGraw Hill, 1990.
- [29] V. Milutinovic, Ed., Microprocessor Design for GaAs Technology, Prentice Hall Advanced Reference Series in Engineering, 1990.
- [30] M. Katevenis, Reduced Instruction Set Computer Architectures, MIT Press, 1984.
- [31] J. R. Ellis, Bulldog: A compiler for VLIW Architectures, MIT Press, 1985.
- [32] S. S. Sapatnekar, and S.-M. Kang, Design Automation for Timing Driven Layout Synthesis, Kluwer Academic Publishers, 1993.
- [33] R. Jain, The Art of Computer Systems Performance Analysis, J. Wiley & Sons, 1991.
- [34] S. A. Przybylski, Cache and Memory Hierarchy Design, Morgan Kaufman Publishers, 1990.
- [17] H. B. Bakoglu, Circuits, Interconnections, and Packaging for VLSI, Addison Wesley Publishers, Inc., 1990.
- [18] D. A. Patterson, & J. L. Hennessy, Computer Organization & Design - The Hardware Software Interface, Morgan Kaufman Publishers, 1994.
- [19] E. J. Rymaszewski, Handbook of Microelectronics Packaging, Van Nostrand, 1990.
- [20] F. E. Gardiol, Lossy Transmission Lines, Artech House, 1987.
- [21]. Y. Taur, D. A. Buchanan, W. Chen, D. L. J. Frank, K. Ismail, S.-H. Lo, G. A. Sai-Halasz, R. G. Vishwanathan, H.-J. C. Wann, S. J. Wind, and H.-S. Wong, "CMOS Scaling into the Nanometer Regime, Proc. of the I.E.E.E., Vol. 85(#4), pp. 486-504, April 1997.
- [21.] Y. Taur, C H. Wann, "0.25 nm CMOS Design Considerations," IEDM 1998, pp. 798-800.
- [22.] G. Freeman, D., Ahlgren, D. R. Greenberg, R. Groves, F. Huang, G. Hugo, B. Jagannathan, S. J. Jeng, J. Johnson, K. Schonenberg, K. Stein, R. Volant, S. Subbanna, "A 0.8 micron 90GHz FT SiGe HBT BiCMOS, ASIC-Compatible, Copper Interconnect Technology for RF and Microwave Applications, IEDM 1999, pp. 40-44.

VI. BIBLIOGRAPHY

- [Agar93] Agarwal, A. and S. D. Pudar, "Column-associative caches: A technique for reducing the miss rate of direct-mapped caches," 20th Annual International Symposium on Computer Architecture ISCA '90, San Diego, Calif., May 16-19. *Computer Architecture News* 21:2 (May), 179-90.

- [Beac88] Beach, W. F. and Austin, T. M. "Parylene as dielectric for the next generation of high density circuits," proceedings of the 2nd International SAMPLE Electronics Conference, June 14-16, 1988 pp 25-45.

- [Bens95] Benschneider, Bradley J., A. J. Black, W. J. Bowhill, S. M. Britton, D. E. Dever, et. al., "A 300-MHz 64-b quad-issue CMOS RISC microprocessor," *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 11, Nov. 1995, pp. 1203-1214.

- [Casc91] Cascade Microtech, Incorporated. "Multicontact high-speed integrated circuit probes." Beaverton, Oregon, 1991.

- [Dabr93] S. Dabral, X. Zhang, X. M. Wu, G.-R. Yang, L. You, H. Bakhru, R. Olson, J. A. Moore, T.-M. Lu. and J. F. McDonald, "aa'a'a" Poly-tetrafluoro-p-xylene as an interlayer dielectric for thin film multichip modules and integrated circuits," *Journal of Vacuum Science and Technology*, B 11(5), Sep/Oct 1993.

- [Deve91] Devore, Jay S. *Probability and Statistics for Engineering and the Sciences*, Third Edition. Pacific Grove, California. Brooks / Cole Publishing, 1991.

- [Dill88] Dillinger T. E. *VLSI Engineering*. pp. 624-93, Englewood Cliffs, New Jersey: Prentice Hall, 1988.
- [Faus95] Faust, Bruce. "Designing Alpha-based systems." *Byte Magazine*, pp. 239-240, June 1995
- [GE95] G.E. Corporate Research & Development Advanced Electronics Assemblies Program, "Microwave High Density Interconnect Design Guide." February 1995
- [Greu90] Greub, H. J. "FRISC - A fast reduced instruction set computer for implementation with advanced bipolar and hybrid wafer scale technology." Ph.D. dissertation, Rensselaer Polytechnic Institute, Troy, New York, December 1990.
- [Greu91] Greub, H. J., et. al. "High-performance standard cell library and modeling technique for differential advanced bipolar current tree logic." *IEEE Journal of Solid-State Circuits*, Vol. 26, No. 5, pp. 749-62, May 1991.
- [Hall93] Haller, T. R., et. al. "High frequency performance of GE high density interconnect modules." *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. 16, No. 1, pp. 21-27, February 1993.
- [Henn96] Hennessy, J. L., and D. A. Patterson. *Computer Architecture: A Quantitative Approach, second edition*. San Mateo, California: Morgan Kaufmann, 1996.
- [Hill84] Hill, Mark D. and Alan Jay Smith. "Experimental evaluation of on-chip microprocessor cache memories," *Proc. Eleventh International Symposium on Computer Architecture*, June 1984, Ann Arbor, MI, 1984.
- [Kilb62] Kilburn, T., D. B. G. Edwards, M. J. Lanigan, and F. H. Sumner. "One-Level Storage System." *IRE Transactions on Electronic Computers*, Vol. EC-11, No. 2, pp. 223-236, April 1962.

- [Lev95] Lev, Lavi A., A. Charnas. M. Tremblay, A. R. Dalal, B. A. Frederick, et. al., "A 64-b microprocessor with multimedia support," *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 11, Nov. 1995, pp. 1227-1236/
- [Long90] Long, S. I., S. E. Butner. *Gallium Arsenide Digital Integrated Circuit Design*, New York, McGraw-Hill Publishing Company, 1990.
- [Maie94] Maier, C. "A testing scheme for a sub-nanosecond access time static RAM" Masters Thesis, Rensselaer Polytechnic Institute, 1994.
- [Maji89] Majid, N., Dabral, S., and J. F. McDonald. "The parylene-aluminum multilayer interconnection system for wafer scale integration and wafer scale hybrid packaging." *Journal of Electronic Materials*, Vol. 18, No.2, pp. 301-311, 1989.
- [Matt70] Mattson, R. L., J. Gecsei, D. R. Slutz, and I. L. Traiger. "Evaluation techniques for storage hierarchies." *IBM Systems Journal*, 9, pp. 78-117, 1970.
- [Maun86] Maunder, C. "Paving the way for testability standards." *IEEE Design and Test of Computers*, Vol. 3, No. 4, p. 65, 1986.
- [Maun92] Maunder, C. M. and R. E. Tulloss. "Testability on TAP." *IEEE Spectrum*, pp. 34-37, February 1992.
- [Nah91] Nah, K., R. Philhower, J. S. Van Etten, S. Simmons, V. Tsinker, J. Loy, H. Greub, and J. F. McDonald. "F-RISC/G: AlGaAs/GaAs HBT standard cell library," *Proc. 1991 IEEE International Conference on Computer Design: VLSI In Computers & Processors*, pp. 297-300, 1991.
- [Nah94] Nah, K. "An adaptive clock deskew scheme and a 500 ps 32 by 8 bit register file for a high speed digital system" Ph.D. Dissertation, Rensselaer Polytechnic Institute, 1994.

- [Phil93] Philhower, B. "Spartan RISC architecture for yield-limited technologies" Ph.D. Dissertation, Rensselaer Polytechnic Institute, 1993.
- [Przy90] Przybylski, S. A. *Cache and Memory Hierarchy Design: A Performance-Directed Approach*. San Mateo, California: Morgan Kaufmann, 1990.
- [Salm93] Salmon, Linton G. "Evaluation of thin film MCM materials for high-speed applications." *IEEE Trans. On Components, Hybrids, and Manufacturing Technology*, Vol. 16, No. 4, June 1993.
- [Ston90] Stone, Harold S. *High Performance Computer Architecture*, Second Edition. Reading, Massachusetts. Addison-Wesley, 1990.
- [Sze81] Sze, S. M. *Physics of Semiconductor Devices*. Second Edition, pp. 182-3, New York: John Wiley and Sons, 1981.
- [Sze90] Sze, S. M. *High-Speed Semiconductor Devices*. pp 371-373, New York: John Wiley and Sons, 1990.
- [Tien95] Tien, C.-K. "System design analysis, implementation, and testing of a 32-bit GaAs microprocessor" Doctoral Thesis, Rensselaer Polytechnic Institute, 1995.
- [Webe92] Weber, S. "JTAG finally becomes an off-the-shelf solution." *Electronics*, Vol. 65, No. 9, p. 13, 10 August 1992.
- [Zhan95] Xin Zhang, "Parylene as an interlayer dielectric," Ph.D. Dissertation, Rensselaer Polytechnic Institute, 1995.

VII. APPENDIX

Optimization of the Register File and Cache RAM Blocks Used in the RPI Datapath and Cache RAM Chips

From the HSCD test structures, we were able to glean more information about the Rockwell interconnect and device performance. This information was backannotated into our models and CAD tools in order to predict the performance of our circuits. What we found was that the memory circuits would perform significantly below their required levels (the target performance for the register file was 200 ps and 450 ps for the cache RAM block). This is due in part to increased parasitic capacitance as well as degraded device performance. Figure 88 shows a comparison of the register file access times for different interconnect and device models. The latest device models are called "2-sided" and "3-sided" and the most recent interconnect model is called "anisotropic".




Figure 88. Register file performance with various device and interconnect models.

Because the register file had tighter performance requirements than the cache RAM blocks, it was selected for optimization first. Although any layout improvements from the register file could be fed into the cache RAM block, separate circuit optimizations were required due to the differences in size and power. For the most part, the cache RAM block optimization followed the same process as used in the register file and was somewhat easier due to the larger access time requirements. To date, both the register file and cache RAM blocks have been redesigned to meet their performance

requirements. The register file has a 195 ps READ access time and the cache RAM block has a 400 ps READ access time. These performance metrics (along with some safety margin) have been incorporated into the redesign of the datapath and cache RAM chips.

Register File / Cache RAM Optimization Process

The optimization process began with layout because the process design rules had changed but the physical design had not been updated. There are three sets of large nodes in the register file and cache RAM blocks, namely the address lines, bit lines, and word lines. The capacitance of these nodes has a direct effect upon performance. For, it can be seen that the largest contribution to the access time comes from the memory cells and bitlines, followed by the address drivers and address lines, and finally the word drivers and word lines.

Although the relative contributions to delay were known, the effect of layout optimizations upon each delay component was not. A series of simulations using SPICE were performed in which the capacitance of the address, bit and word lines was varied in order to determine the sensitivity of the circuit delay to that component. The results (shown in Figure 88) indicated that the bit and word lines are the most sensitive, suggesting that the optimization process should focus upon these nodes. Due to the relatively large bit line capacitance (~3X larger than the word line value), the bit lines became the primary focus of layout optimization.




Figure 89. Access time sensitivity to address, bit and word line capacitance.

Layout Optimizations

Now that the sensitivity of the access time to the node capacitances was known, the emphasis shifted to minimizing capacitance through layout changes. The primary focus was the bit lines in the memory cells. The word lines were also optimized indirectly as a side-effect of the bit line changes. Although the register file had much lower sensitivity to the address lines, they were optimized anyway in order to squeeze out as much performance as possible.

Bit / Word Line Optimizations

A number of memory cell layouts have been progressively developed, some of which have been made possible by process design rule changes. To date eleven distinct memory cell layouts (**Error! Reference source not found.**) have been produced along with numerous variations. The first four iterations produced the most significant improvements in parasitic capacitance but unfortunately they were not sufficient alone in meeting the target performance numbers. Circuit modifications were then undertaken and a new memory cell was developed (described in the next section).

The original memory cell had several disadvantages, which were solved with the addition of metal-3 to the Rockwell process. The primary problem was the parasitic capacitance between the metal-1 bit and the metal-2 word lines. The first iteration placed the top word line in metal-3 to reduce the crossover capacitance. This helped somewhat but it wasn't sufficient. The justification for leaving the lower word line in metal-2 was to avoid the large metal-2/metal-3 via which would be required to connect a metal-3 word line to the metal-1 resistor connection. For the upper word line, this via could be hidden underneath the resistors, but for the lower word line a via would complicate routing and possibly increase the coupling between the lower word line of one row and the upper word line of the next. In the end, it was decided that routing the lower word line in metal-3 was necessary despite the disadvantages, so the second cell iteration was produced.

The next redesign opportunity arose when Rockwell reduced the dimensions of the HBT devices and relaxed the minimum feature sizes. These changes allowed the memory cell to be packed more tightly, creating more room for the bit lines and reducing their parasitic capacitance. The smaller feature sizes also allowed the resistors to be shrunk which became important in later redesigns. The effects of the process/design changes can be clearly seen in the fourth iteration of the memory cell: the resistors and devices are smaller, the devices are placed closer together and the interconnect is routed closer to the devices. Since the core of the cell is now more compact than before, the coupling to the bit lines is reduced because the adjacent structures are further away. More importantly, a smaller core allows the bitline - bitline spacing to be increased. Because the majority of the bitline coupling is with the neighboring bitline, any reduction can significantly improve the overall bit line parasitic capacitance. After the core is redesigned for maximum compactness, the bitline-bitline spacing is adjusted to determine the optimum spacing for minimal parasitic capacitance.




Figure 90. Memory cell layouts.

Register File Circuit Sensitivity Analysis and Component Modifications

Once it became apparent that layout modifications alone would not be sufficient to recover the "lost" performance, a series of SPICE simulations were performed to determine the sensitivity of the register file to component value changes. There are numerous components in the register file, which have an impact upon the performance, but several components are particularly important. These are the address decoder resistors, the read/write logic pull-ups and current source resistors, the sense amplifier bitline current source resistor, and the memory cell resistor ratio in the threshold voltage generator. Some components affect several nodes in the circuit with conflicting requirements, presenting a difficult and complex optimization problem.

Address Decoder: Wordline Voltage

The address decoder directly sets both the address line and wordline voltage swings. The wordline is the mechanism by which a row of memory cells is selected and enabled to place their logical values on the bitlines. As a result, the switching time of the wordlines directly impacts the overall access time of the register file. The wordline swing is

determined by both the total resistance in the address decoder, the ratio of the resistors and the V_{BE} of the devices.

In Figure 91, the effect of different total decoder resistance values upon the access time and wordline swing are shown for a top:bottom resistor ratio of 1:1. When the total resistance is increased, the wordline swing also grows because the voltage drop across the total resistance increases, forcing the wordline driver base lower and thus the wordline voltage as well. The upper value of the wordline voltage is fixed at $V_{CC} - V_{BE}$ because the base is pulled to V_{CC} when all five of the address decoder Q1s are cut-off. From the second plot in Figure 91, a lower bound on the total resistance of about 420 Ω can be determined which will satisfy the minimum wordline static swing of 850 mV.




Figure 91. Wordline sensitivity to total decoder resistance.

Figure 92 depicts the effect of various decoder resistor ratios for a total resistance of 440 Ω . Although it appears that the wordline swing should not be affected by the resistance ratio, the changing ratio does reduce the current through the Q1 devices. The different current levels in turn affect the voltage drop across the total decoder resistance and thus

the wordline swing. This is just one example of the intricate and complex balance between different parts of the register file circuit.




Figure 92. Wordline and address-line current sensitivity to decoder resistor ratio (total resistance = 440 Ohms).

Address Decoder: Address Line Voltage and Current

The address lines are also affected by the resistors in the address decoder. The ratio determines the voltage swing on the address lines, which in turn determines the current. The maximum address line voltage is fixed at $V_{CC} - V_{BE}$ but the decoder ratio determines the minimum voltage and thus the total swing. Current flows through the address lines only when the address line is low, hence the current decreases with increasing swing (or, alternatively, the current decreases with decreasing minimum address line voltage).




Figure 93. Address line sensitivity to decoder resistance ratio.

Read/Write Logic: Bitline Voltage

The read/write logic has a significant effect upon the bitlines, primarily in the WRITE mode. In order to overwrite the state of a memory cell, the read/write logic pushes the bitlines to relatively extreme high and low voltages in order to cut-off and turn-on the memory cell devices. The speed of the WRITE as well as the recovery time are determined primarily by the bitline swing (a larger range results in faster WRITES but a slower recovery and vice versa). During a READ, the circuit attempts to set the bitlines to a mid-range value. This specifies the low bitline voltage and thus clamps the lower part of the bitline swing.

The read/write logic uses the threshold voltage along with resistive pull-ups and a resistor current source to generate the bitline voltages. The actual bitline potentials depend upon the value of the pull-up resistors and the amount of current flowing through the resistors (determined by the resistive current source). For a READ, current flows through both resistors equally, dropping the resistance by half and producing a mid-range voltage of $V_{th} - IR$. During a WRITE, current only flows through one of the resistors, hence the voltage swing is V_{th} to $V_{th} - IR$. Because the read/write logic uses the threshold voltage V_{th} as a reference and power supply, drawing excessive amounts of current from the threshold voltage generator can seriously stress the generator circuit and reduce its robustness. For this reason, the amount of current, which can be drawn by the read/write logic, is limited and should be kept low if possible.

Figure 94 below shows the time required to perform a WRITE and the bitline swing for a range of read/write logic current source resistances (also shown are the static READ access times). As the current source resistance increases, the current through the pull-ups decreases and the bitline swings are reduced. This leads to longer WRITES and eventually (at higher resistance values) failure to overwrite the memory cell state.




Figure 94. Read/write logic bitline swings during WRITE
(pull-up resistance = 600 Ω).

The pull-ups in the read/write logic have an even larger effect upon the WRITE time but can adversely affect the READ times by lowering the low bitline voltage and increasing the total bitline swing. Most importantly, a larger pull-up resistor increases the time required to switch from a WRITE to READ because the internal read/write logic swings are higher and thus more charge must be dissipated to change modes. In the end, however, the choice for the current source resistance was made to reduce the strain on the threshold voltage generator and the pull-up value was optimized for this current. Figure 95 below shows the effect of different pull-up values on the access times and the bitline swing during a WRITE.




Figure 95. Access times and bitline swing during WRITE for various read/write logic pull-up resistor values.

Sense Amplifier: Bitline Current

The sense amplifiers contain the current source for the bitlines. By varying the current through the bitlines, the delay due to parasitic capacitance can be significantly reduced. However, care must be taken not to burn out the devices in the memory cells, hence the maximum bitline current is limited.

The bitline current source is simply a high-current device and a resistor connected between the emitter and V_{EE} . A bias generator sets the base voltage and produces a constant voltage drop across the tail resistor, thereby determining the bitline current. Because the bitlines exhibit the most sensitivity to capacitance of all the large nets in the register file, they offer the most opportunity for improvement. By increasing the current flowing through the bitlines, they can be discharged quickly and thus improve the switching time. Figure 96 demonstrates the sensitivity of the register file access time to




Figure 96. Access time sensitivity to bitline current.

the bitline current.

Register File Circuit Modifications


After the sensitivity analysis was performed and the circuit components were fine-tuned, it was obvious that component value changes alone would also not be sufficient to meet the performance requirements. In order to correct some of these problems and improve performance, various circuit modifications were explored and analyzed. The optimization process focused upon the static and dynamic circuit performance in order to

reduce the static access time. The difference between static and dynamic performance was primarily felt at two locations in the circuit: the wordlines and bitlines.

Wordline Swing / Memory Cell

Because the wordlines provide the means for selecting a row within the register file and determine the bitline swing, their switching time has a direct impact upon the performance of the register file. The layout modifications have trimmed the parasitic capacitance down tremendously and the resistance of the line is minuscule, hence the RC effect does not contribute significantly to the delay. However, the device switching speed does contribute and is dominant, hence the swing of the wordline has a direct effect upon the switching time.

As can be seen in Figure 97(b), the wordline swing for static and dynamic signals are significantly different. Because the static swing is higher, the time required to switch the wordline after it has charged fully (due to a "static" address) is greater than if the address had changed in the previous cycle (i.e. a "dynamic" address). The wordline swing determines the swing on the memory cell collector nodes, which drive the bitlines, thus when the wordline switching time increases, it directly affects the bitline switching time. Ideally, the static and dynamic swings should be equal, eliminating any difference between access times.



(a) Bitline swings (static, dynamic) (b) Wordline swings (static, dynamic, static)

Figure 97. Internal signal swings due to relatively static and dynamic address changes.

Wordline clamp

Several clamping circuits were investigated as a way to restrict the high static wordline swing, but severe operating requirements hampered this effort. Some of the problems were the high wordline current (approximately 20 mA), the large 0.8 V drop across the Schottky diodes, and the need to fit the clamp circuit within a small area in order to

maintain the original register file dimensions. In the end, no satisfactory circuit was found.

Wordline voltage divider

Because the switching time is actually based upon when the bitlines switch rather than the wordlines, there is the possibility of improving the access time without reducing or limiting the wordline swing. By lowering the internal memory cell swings, the bitline swing is also reduced and thus switches faster when the wordlines start to change.

One drawback to reducing the internal swings was the reduction of the dynamic wordline swing and a corresponding increase in the dynamic access time. However, although the dynamic access time increases significantly, it is still less than the static access time. Since only the longest access time is important from the standpoint of the F-RISC/G datapath chip, the relatively fast dynamic access time of the original design provided no benefit and could be sacrificed for the benefit of the static access time.

To reduce the internal memory cell swings, a simple voltage divider was created in the memory cells by placing a resistor between the wordline and the previous wordline connection point (see Figure 98). This resistor provides a voltage drop and creates an "effective" wordline. The actual potential drop across the resistor depends upon the selected/deselected state of the memory cell due to the different current levels. Because the drop is proportional to the current, it reduces the effective wordline potential in the selected state much more than in the deselected state.

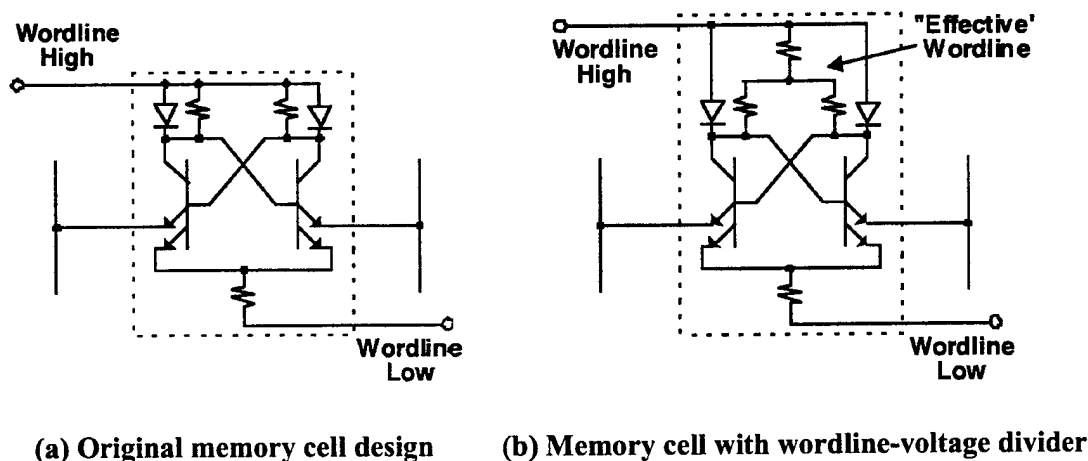


Figure 98. Original and reduced-wordline-swing memory cell designs.

Simulations in SPICE agree with the analysis above. The static access times are reduced at the expense of the dynamic access times. The addition of a small voltage-divider

resistor into the memory cells was simpler to implement than a clamping circuit for each row in the register file and did not increase device count, further justifying this method.

Bitline Swing: Read/Write Logic

During a READ, only the high bitline is actually driven by the memory cells. The low bitline voltage is set by the read/write logic based upon the threshold voltage. For a WRITE, the read/write logic sets both bitline voltages in order to overwrite the memory cell state. To do this, the read/write logic has to force the bitlines to values, which will force the devices in the memory cell on or off and thereby store the logical value. The bitline voltages during a WRITE determine in part the speed of the operation with larger bitline swings corresponding to faster WRITES. However, after the WRITE operation is over, a new address may be presented for a READ and the register file must respond with the appropriate data within 200 ps. If the WRITE bitline voltages are too large, the switching time of the bitlines may be significantly delayed due to the excess charge from the WRITE. One way to avoid this situation is to increase the lower bitline voltage while decreasing the high one.

The original design of the read/write logic applied high and low bitline voltages of equal magnitude relative to the voltage of the read/write logic during a READ. A sensitivity analysis was performed in which the magnitude of the bitline swing during a WRITE was varied and the time to store the data was measured. The results indicated that the WRITE time was within the specifications while the bitline swing was below the normal READ levels, meaning that the swing during a WRITE did not have to be adjusted.

Even though it was not required in the register file, adjusting the bitline swings during a WRITE was necessary in the cache RAM optimization. During the redesign of the register file, it was not clear that no changes were necessary regarding the WRITE bitline voltage swings and a new read/write logic circuit was developed which reduced the high bitline excursion. The read/write logic operates by generating three distinct voltages: a mid-range voltage for both bitlines during a READ and high and low voltages for the bitlines during a WRITE. All of the voltages are based upon the threshold voltage and the mid-range and low voltages are generated using resistors.

Bitline Swing: Bridge Resistor

To improve the switching performance of the bitlines, a "bridge" resistor was connected between them (Figure 99). The bridge resistor attempts to equalize the bitline voltages (and thereby improve the switching speed) but is large enough to maintain the bitline swing between address changes. The actual value of the bridge resistor was determined using a sensitivity analysis, which examined the access time, WRITE time, bitline swing, memory cell device current and current through the bridge resistor.

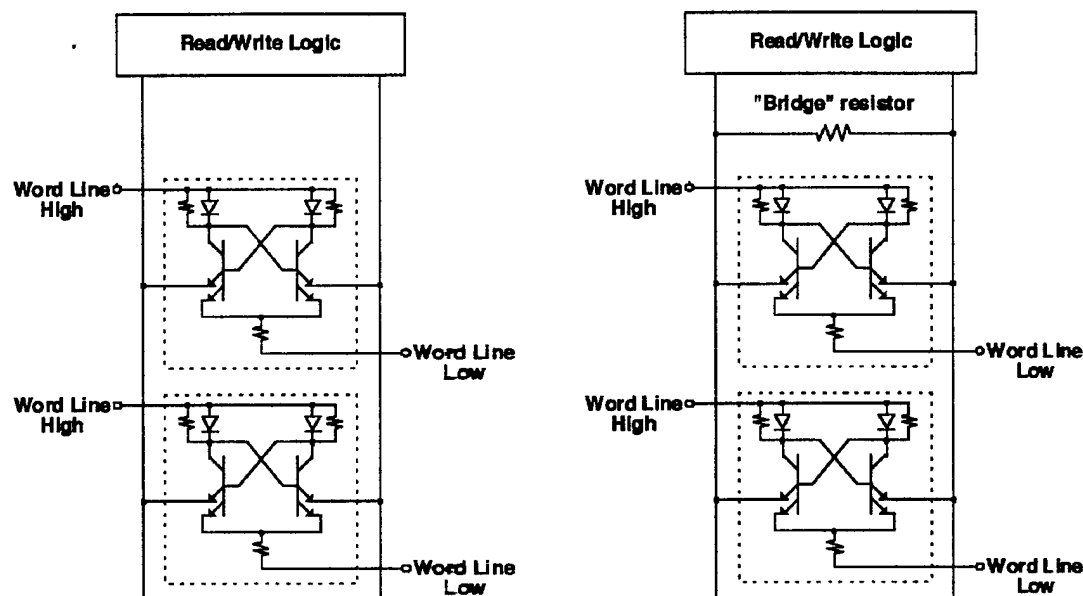


Figure 99. "Bridge" resistor between bitlines.

The bridge resistor affected many parts of the register file circuit. It increased the current through the memory cell devices significantly because, in addition to sinking current from the bitline current sources, current was also coming from the other bitline. It also increased the WRITE time for the same reason but to a greater extent due to the larger bitline swings during the WRITE. Despite all of these negatives, the bridge resistor increased the register file access time significantly. Figure 100 shows the effect of various bridge resistor values upon the bitline swing, the bitline current and the current through the resistor itself. In Figure 101, the static and dynamic READ access times and the WRITE time are shown relative to the bridge resistor value.





Figure 100. Sensitivities to bitline bridge resistor.




Figure 101. Performance sensitivity to bitline bridge resistor.

Register File Optimization Summary

The 32x8 register file circuit and layout has been optimized to achieve a 195 ps READ access time using 2.01 W. The 32x16 cache RAM block has a 400 ps READ access time with a power dissipation of 1.5 W. The external dimensions of the register file have remained the same while the cache RAM blocks were increased by 7 μm to accommodate the bridge resistor (the cache RAM block requires a 3k Ω resistor, significantly larger than the register file resistor).

VIII. SAMPLE PUBLICATIONS

High-Performance Standard Cell Library and Modeling Technique for Differential Advanced Bipolar Current Tree Logic

Hans J. Greub, *Member, IEEE*, John F. McDonald, *Member, IEEE*, Ted Creedon,
and Tadanori Yamaguchi

Abstract—A high-performance standard cell library for the Tektronix advanced bipolar process GST1 has been developed. The library is targeted for the 250-MIPS Fast Reduced Instruction Set Computer (FRISC) project. The GST1 devices have a minimal emitter size of $0.6\text{ }\mu\text{m} \times 2.4\text{ }\mu\text{m}$ and a maximum f_t of 15.5 GHz. By combining advanced bipolar technology and high-speed differential logic, gate propagation delays of 90 ps can be achieved at a power dissipation of 10 mW. The fastest buffers/inverters have a propagation delay of only 68 ps. A 32-b ALU partitioned into four slices can perform an addition in 3 ns using differential standard cells with improved emitter-follower outputs and fast differential I/O drivers. A modeling technique for high-speed differential current tree logic is introduced. The technique gives accurate timing information and models the transient behavior of current trees.

I. INTRODUCTION

THIS PAPER describes an experimental standard cell library for the advanced bipolar process GST1 under development at Tektronix [1]. The cell library was designed for the Fast Reduced Instruction Set Computer (FRISC) project [2], [3]. The 32-b processor is partitioned into circuits with a maximum complexity of 1000 current tree gates since the process yield is too low for a single-chip implementation. The standard cell library was optimized for speed to achieve a processor cycle time of 4 ns. Differential ECL logic is used to lower propagation, interconnect delays, and switching noise.

Because of the low targeted complexity, a higher power dissipation could be accepted in the speed versus power trade-off than in previously reported advanced bipolar libraries [4], [5]. The current trees are built out of the smallest GST1 devices with ECL output drivers to lower interconnect loading delays. The resulting standard cells are characterized by a high drive capability combined with a low fan-in load. Each cell is made available with three

different output drivers. The drive capability of the ECL output driver circuits was improved to reduce the need for high-power gates.

The transient behavior of the high-speed, high-power cells is not dominated by interconnect capacitance as current starved ECL. Thus transients and glitches intrinsic to the structure of current trees are visible and cannot be simulated with a simple behavioral logic model. A structural modeling technique for high-speed current tree logic has been developed to improve the delay accuracy and to capture transients that could lead to circuit failure.

To reduce I/O delays, high-speed differential drivers and receivers with a low logic swing of $\pm 250\text{ mV}$ are provided besides standard single-ended I/O circuits. The single-ended drivers are ECL 10K compatible and have a voltage swing of 865 mV. Low I/O delays are crucial for the carry propagation in the FRISC data path, which had to be partitioned into four 8-b slices. In particular, the 32-b ALU is on the most critical delay path of the processor and is, therefore, examined in detail.

II. ADVANCED BIPOLAR CIRCUIT TECHNOLOGY

A. Advanced Bipolar Process

The GST1 advanced bipolar n-p-n transistor devices are built with a self-aligned polysilicon emitter-base (E-B) process with a coupling base implant. This results in shallow emitter and base junction depths [6]. Fig. 1 shows the structure of n-p-n devices and polysilicon resistors and Fig. 2 shows a SEM device cross section. The $1\text{-}\mu\text{m}$ trench isolation reduces the collector-to-substrate (C-S) capacitance and increases device density. The smallest devices have an emitter stripe of $0.6\text{ }\mu\text{m} \times 2.4\text{ }\mu\text{m}$ and can be placed on a dense $8\text{-}\mu\text{m} \times 12\text{-}\mu\text{m}$ grid. A self-aligned titanium-silicide layer on top of the polysilicon layer for emitter and collector contacts reduces the sheet resistance to $1\text{ }\Omega/\square$ and thereby provides an additional layer for short interconnect. The same polysilicon layer without the silicide is used for resistors. Two gold metal layers with a $4\text{-}\mu\text{m}$ pitch are available for interconnect. The advanced bipolar n-p-n devices have a maximum f_t of 15.5 GHz [1]. Ring-oscillator delays of 55 ps per stage have been measured. Further, dual 4-b analog-to-digital converters with a performance of 1.5

Manuscript received September 20, 1990; revised January 15, 1991. This work was supported in part by DARPA under Contract DAAL03-90-G-0187.

H. J. Greub and J. F. McDonald are with the Center for Integrated Electronics, Rensselaer Polytechnic Institute, Troy, NY 12181.

T. Creedon was with the Electronic Systems Laboratory, Tektronix, Inc., Beaverton, OR 97007. He is now with Kestrel Technologies, Lake Oswego, OR 97035.

T. Yamaguchi is with Integrated Circuit Operation, Tektronix, Inc., Beaverton, OR 97007.

IEEE Log Number 9143201.




Fig. 1. Device structure.

Gs/s have been demonstrated [7]. The dimensions and key parameters of the smallest GST1 device are summarized in Table I.

B. Current Switch

Fig. 3 shows the basic building block of current tree logic, the current switch (CSW). The input current into the common-emitter node is switched left or right depending upon the two base voltages. Using a simplified Ebers-Moll model for the bipolar transistor, the dc characteristics of a current switch buffer can be expressed in a closed form [8]. However, the effect of the parasitic emitter and base resistances should be included to obtain a good match. Unfortunately, the analysis does not yield a closed-form solution even if only the emitter resistance R_e is included.

Fig. 4 shows the delay of current switch buffers with and without 500 μm of interconnect capacitance as a function of switching current. The switching current was fixed at 400 μA since increasing the switching current any further would mainly increase power dissipation. The nominal switching current should be set below the optimal current to avoid operation in the region where delays increase rapidly with higher current and power. With a logic swing of ± 250 mV and a fan-out of 1, a current-mode logic (CML) buffer has a delay of 64 ps and an emitter-coupled logic (ECL) buffer with an emitter-follower current I_{ef} of 800 μA has a delay of 66 ps. The CML buffer dissipates only 2 mW but has a propagation delay sensitivity R_s of 400 Ω . The delay sensitivity R_s multiplied by the load capacitance C_l gives the incremental gate delay due to interconnect loading. A linear delay dependence is a good approximation for ECL or CML circuits [9]:

$$T_d = T_0 + \Delta t_d = T_0 + R_s \cdot C_l.$$

The ECL buffer has a power dissipation of 10 mW with an R_s of only 119 Ω . To save power, CML is used within standard cells where the interconnect length is short. The nominal voltage swing was fixed at ± 250 mV. This drives the current switch well beyond the points with maximum noise margin (gain = 1) and results in a voltage gain of 2.6 at 360 K. The voltage swing is determined by a trade-off between the delay sensitivity to capacitive loading and the desired voltage gain and noise margin of the logic. The



Fig. 2. SEM cross section of n-p-n device.

TABLE I
GST1 MINIMAL N-P-N DEVICE PARAMETERS

Size	8 $\mu\text{m} \times 12 \mu\text{m}$
Emitter Size	0.6 $\mu\text{m} \times 2.4 \mu\text{m}$
Current Gain h_{FE}	100
Emitter Resistance R_e	60 Ω
E-B Capacitance	6.7 fF
B-C Capacitance	7.5 fF
C-S Capacitance	9.0 fF
CutOff Frequency f_t , $V_{CB} = 0.85$ V, $T = 300$ K	12 GHz

current must be fully switched left or right at nominal input voltage levels, otherwise logic level degradation will occur if current switches are cascaded or stacked to build current trees. Fig. 5 shows the buffer delays as a function of logic swing V_l . The delays of CML buffers increase rapidly at high logic swings because the devices start to saturate.

The voltage swing is an important characteristic of a logic family since propagation and interconnect delays as well as switching noise increase with V_l . The ECL buffer delays with 500 μm of interconnect increase by 17% if the logic swing is changed from differential (± 250 mV) to single-ended levels (500 mV). If the interconnect capacitance C_l is large, the incremental gate delay Δt_d as a function of the logic swing V_l can be approximated by

$$\Delta t_d = K_i \cdot \frac{V_l \cdot C_l}{I_s} = K_i \cdot R_l \cdot C_l.$$

The constant K_i can be derived from a sensitivity analysis and depends upon the circuit configuration (ECL = 0.31 for $I_s = I_{ef}$, CML = 0.65) and the device technology [9]. To lower interconnect loading delays either the logic swing V_l must be lowered or the switching current I_s must be increased. Higher switching current implies, however, higher power dissipation. Hence, a low logic swing is the key to high-speed logic with low power dissipation! The switches must exhibit high gain and generate little switching noise to support low logic swings. Bipolar logic with a logic swing of only 250 mV has a big advantage over CMOS with a logic swing of 3–5 V in this respect.

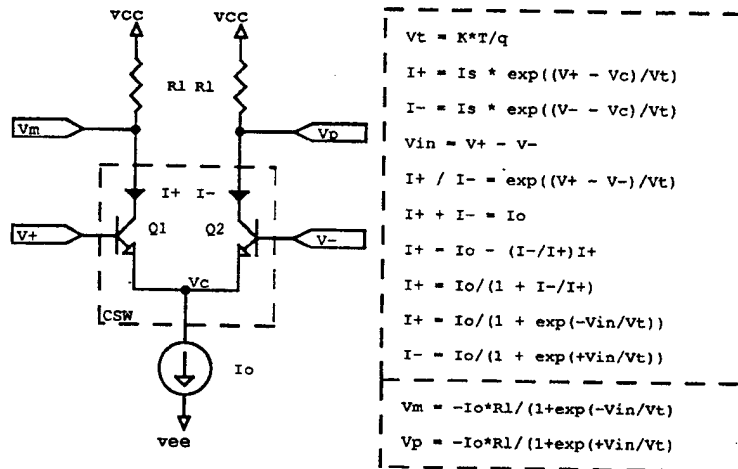


Fig. 3. Bipolar current switch.




Fig. 4. CML and ECL buffer delays versus switching current.




Fig. 5. ECL and CML buffer delays versus logic swing.

Fig. 6 shows the switching noise (power) of differential ($V_i = \pm 250$ mV) and single-ended ECL buffers ($V_i = 500$ mV) for a positive and negative signal transition. The single-ended buffer generates considerable switching noise because of its higher logic swing and unbalanced load. Differential ECL produces only small switching transients and hence evades the delta- I noise problem common to high-speed logic.

C. Differential Current Tree Logic

The high speed and low switching noise of differential logic make it very attractive for bipolar [10], [11] or GaAs logic [12]. Differential GaAs logic is called source-coupled FET logic (SCFL). The high performance and efficient logic implementation of cascaded differential logic trees has led to the development of a similar CMOS logic family at IBM [13], called cascode voltage switch logic (CVSL).

Fig. 7 shows a differential AND/OR gate with three levels of series gating. An equivalent single-ended OR gate needs twice the voltage swing to obtain the same noise margin. Twice the voltage swing is sufficient, despite the fact that the generation of the reference voltages is sensitive to supply voltage drops on power rails, because doubling the voltage swing also doubles the maximum gain of the current switch. To obtain twice the voltage swing either the load resistance R_l or the switching current I_s must be increased by a factor of 2:

$$\text{gain}_{\max} = g_m \cdot R_l = \frac{I_s}{2 \cdot V_T} \cdot R_l = \frac{V_i}{2 \cdot V_T}$$

The number of switches that can be stacked with standard ECL supply voltages is limited to three for ECL and to four for CML. The input signals for current switches at different levels must be offset by at least one base-emitter junction voltage $V_{BE0} \approx 0.85$ V to avoid saturating the bipolar devices. The nominal logic swing at each level is ± 250 mV.

Since a full current tree with three levels of current switches forms a 3-to-8 decoder, any Boolean function of three variables can be implemented in a single current tree by using collector dotting at the top level. An efficient logic implementation is obtained by eliminating current switches with both collectors connected together and by using collector dotting at level two for intermediate decoding states. A four-input multiplexer gate can also be implemented with a single current tree as shown in Fig. 8. By using feedback from the outputs of the current tree, data latches with any two-input gate at the input can be implemented as shown in Fig. 9. The feedback signals are taken from the top of the tree rather than from the output because of layout considerations.

Differential signals can be inverted with zero delay and power by exchanging the true and inverted signal pair connections at any input or output port. This reduces the number of cells in the standard cell library since dual gates like AND/OR are physically identical. Dual gates get mapped into the same cell during netlist generation.

Emitter followers are used to increase the drive capability of the gates and to shift output levels. A standard cell can drive only one output level because emitter followers tend to ring if they have to drive outputs at multiple levels. Since most of the power is dissipated in the emitter followers, each logic gate is available with three different strength drivers ($I_{ef} = 400, 800, \text{ and } 1200 \mu\text{A}$).

The propagation delays of differential logic depend upon the path the current takes through the tree. The delay from inputs at a given level to the top of the tree can, therefore, depend upon input signals at higher levels. For example, in the differential AND gate shown in Fig. 7, the delay from the lowest level input depends upon whether the current flows through current switch S_2 to q or through S_2 and S_3 to q or qb . The maximum propagation delays for a medium-power AND gate with a level-one output are 90 ps from level 1, 135 ps from level 2, and 180 ps from level 3. An equivalent three-input single-ended OR gate has a propagation delay of 95 ps for the OR output. The delay sensitivity R_s of a single-ended OR gate is 131Ω for the rising edge and 257Ω for the falling edge at a power dissipation of 10.5 mW. The medium-power differential AND/OR gate has a power dissipation of 10 mW and a delay sensitivity R_s of only 116Ω . The differential gate has no decisive speed advantage over the single-ended gate at low loads, but the interconnect delay sensitivity of the differential gate is considerably lower and does not depend upon the signal transition.

While differential logic is faster than single-ended logic due to its low logic swing and can be efficiently implemented with current trees, there are also disadvantages. Twice as many signal interconnections must be routed. This increases the average interconnect length since the width of routing channels and feedthroughs doubles. Further, two emitter followers are needed for every gate, which increases power dissipation. However, differential logic requires no power for inverters or reference voltage generators and its sensitivity towards voltage drops on power rails is low.

Existing CAD tools can easily be modified to support three different signal offset levels, differential signal inversion, and checking for input-level violations that cause saturation in standard cells. However, the designer has to assign signals levels avoiding level violations and keeping the propagation delays on critical paths minimal. The standard cell router should support differential wiring. All differential wires should be routed right next to each other to obtain equal loading on differential nets. Parallel routing of differential signals further reduces crosstalk since crosstalk signals will couple almost equally to both wires and thereby produce mainly common-mode noise, which is largely rejected by current switches.

D. Emitter Followers and Buffers

Emitter followers have a tendency to ring, which leads to long settling times. Propagation delays are quite diffi-




Fig. 6. Switching noise of single-ended and differential ECL buffers.

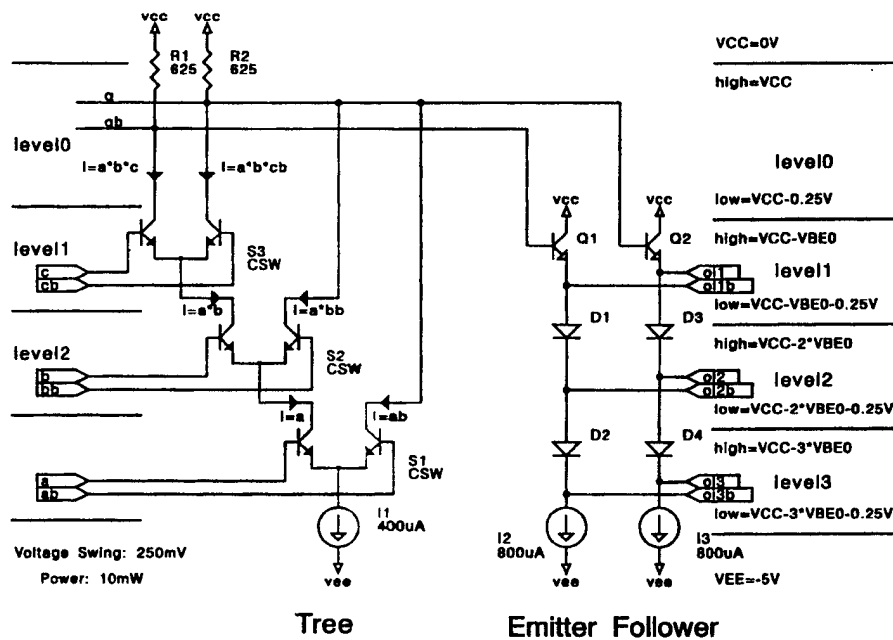


Fig. 7. Differential three-input AND gate.

cult to model for input signals that arrive while the outputs have not yet settled. Therefore, emitter-follower and level-shifter configurations were developed to obtain faster settling times and lower interconnect delays.

The improved emitter followers have a damping resistor between the differential outputs to reduce ringing, as

shown in Fig. 10. For level-2 and -3 emitter followers an f_i doubler circuit is used to reduce ringing and increase driving capability. The damping resistors cause a maximum loss of 20 mV in voltage swing since the current flowing through the emitter-base junction is higher for the transistor with a logic-high output signal. The buffers

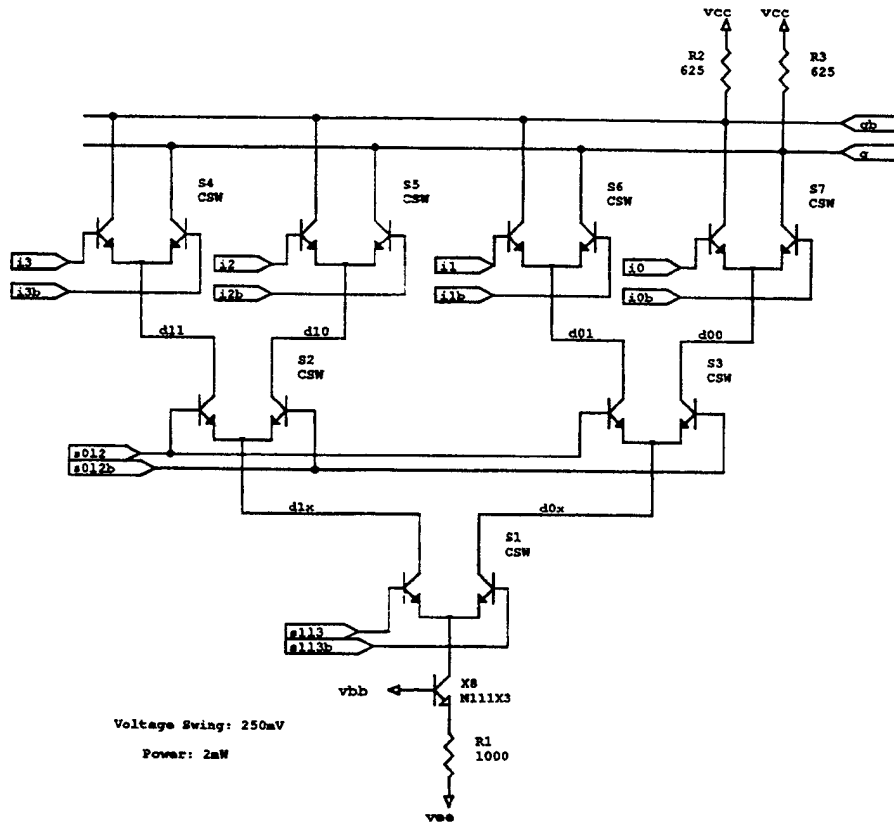


Fig. 8. Differential four-input multiplexor.

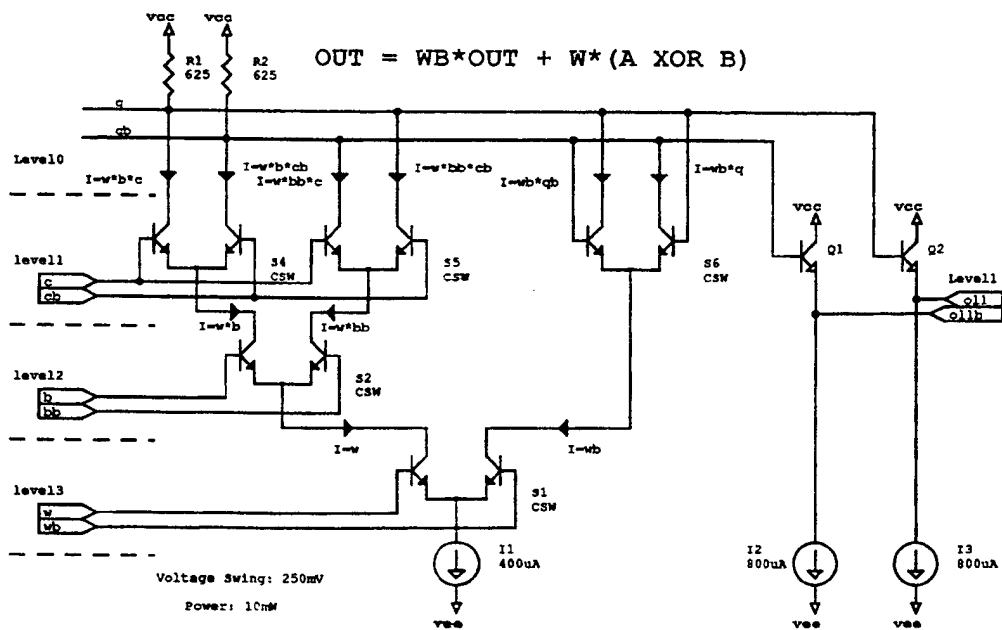


Fig. 9. Differential latch with XOR inputs.

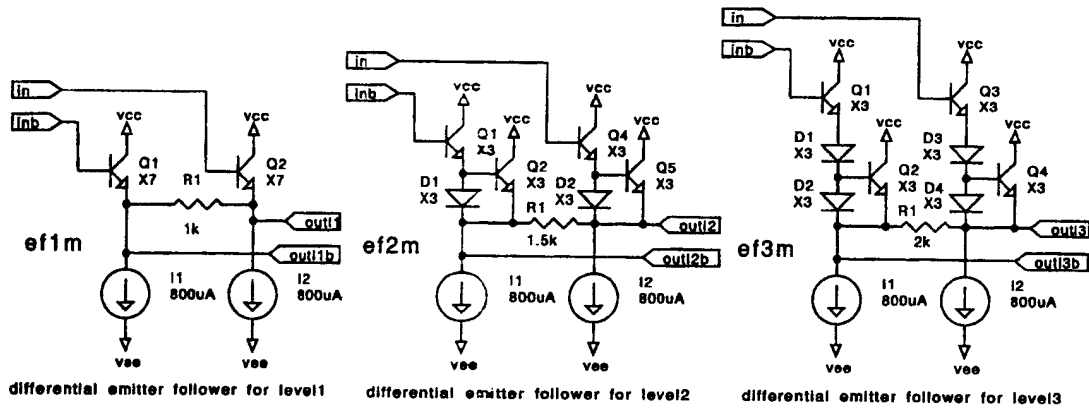


Fig. 10. Improved differential emitter followers.

with the improved differential emitter followers for level 2 and 3 show lower interconnect sensitivities (-21% , -23%). Only the emitter follower for level 1 has an 8 ps higher unloaded propagation delay. However, at high loads the interconnect delays are 17% lower. Without damping resistor the buffer has an underdamped step response with a high overshoot and a long settling time.

Highly loaded emitter followers have largely different rise and fall delays. The rise-time delay is quite small due to the high transconductance g_m of the bipolar devices. The fall time is dominated by the available pull-down current. This leads to highly asymmetrical signal transitions in current-starved ECL.

A special buffer is available for driving long interconnect lines as encountered in clock distribution trees. This super buffer (SBUF1H) has a delay of only 68 ps and a sensitivity R_i of only $60\ \Omega$ at a power dissipation of 12 mW. The SBUF1H circuit shown in Fig. 11 consists of a current switch buffer with a switched current source for the emitter followers. This results in a push-pull output stage with a high pull-down current of 2 mA. Resistor R_3 provides damping and keeps a minimal current of $800\ \mu\text{A}$ flowing through Q_5 or Q_6 . It prevents the high output from slowly charging up to the V_{cc} power level through the base-emitter junctions of the output transistors. The SBUF1H has lower power dissipation and lower interconnect delays than a standard high power buffer, but the fan-in load is three times higher.

E. Input / Output Circuits

High-speed I/O drivers are especially important in advanced bipolar logic since large circuits need to be partitioned because of power dissipation limits and fabrication yields. Two different types of drivers and receivers are provided as shown in Fig. 12. Single-ended 10K ECL-compatible drivers/receivers have a driver plus receiver delay of 300 ps for a rising edge and 312 ps for a falling edge with an I/O pad capacitance of 1 pF. The driver has

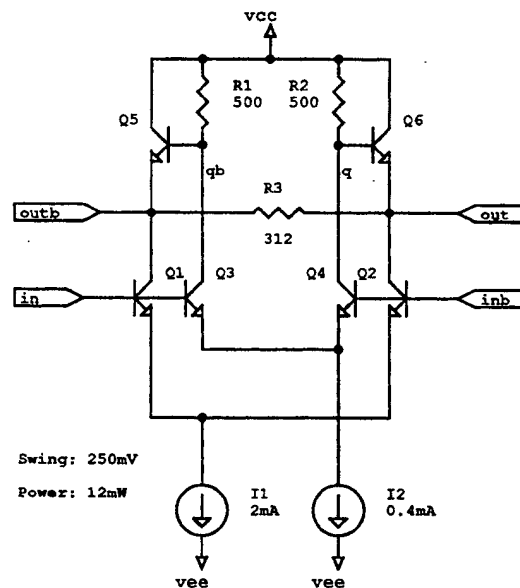


Fig. 11. Super buffer SBUF1H.

the typical unbalanced power dissipation of single-ended drivers. These unbalanced drivers cause considerable delta- I noise because of voltage drops on bondwires and power rails. Therefore, a dedicated power rail V_{pp} (0 V) is used for single-ended drivers to keep the delta- I noise away from the standard cell core.

The high current (16 mA) that is switched on and off by single-ended drivers causes a significant voltage drop on the bondwires, which have an inductance of about 20 pH/mil and are typically 10–15 mils long. Simulations predict 30 mV of delta- I noise for a single-ended I/O driver with a 15-mil bondwire on the V_{pp} power supply. Therefore, only two to three drivers can be supplied with

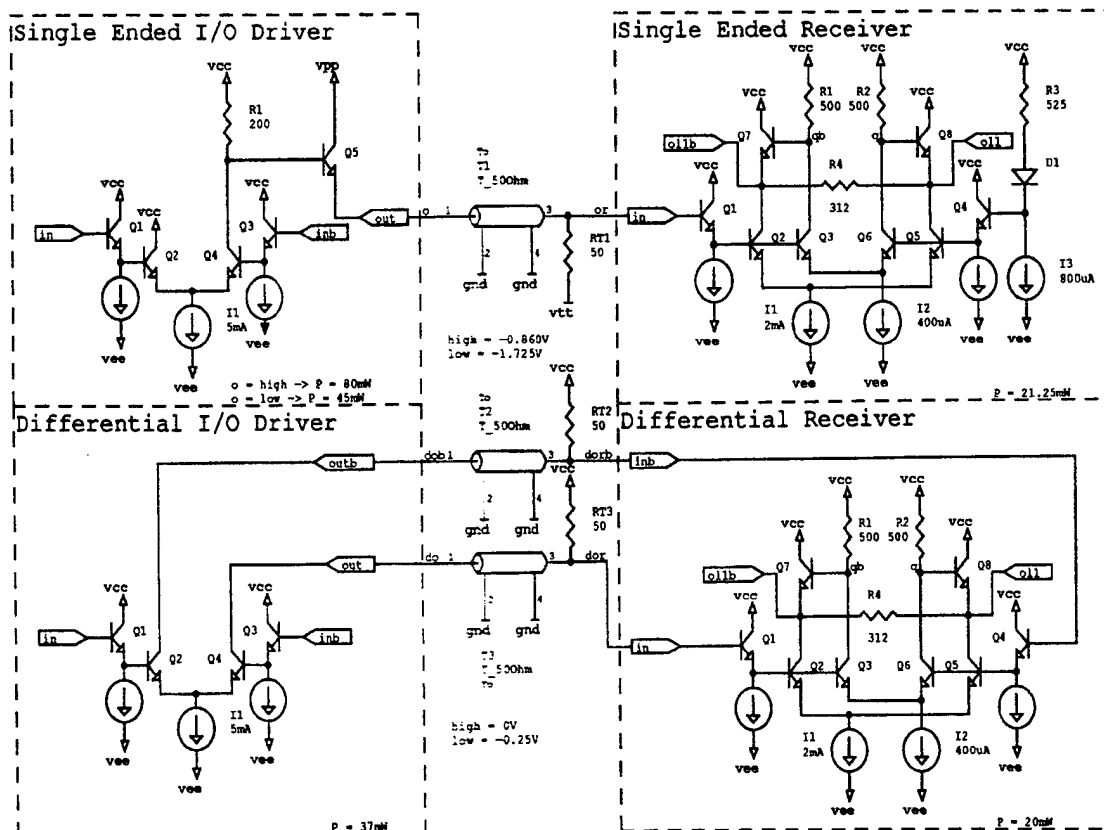


Fig. 12. I/O drivers and receivers.




Fig. 13. Standard cell test circuit.

one V_{pp} power pad else the voltage drop on the bondwire and power rails can cause saturation of the output devices. By using tab bonding or a flip-chip die mount, the power supply inductance could be substantially reduced.

The second driver is a differential open-collector driver with a voltage swing of only ± 250 mV. The two transmission lines are terminated with $50\text{-}\Omega$ resistors to V_{cc} . The differential driver plus receiver delay is only 220 ps with a

pad capacitance of 1 pF. Differential drivers have the disadvantage of using up two I/O pads, however, since they have lower and balanced power dissipation fewer power pads per driver are required. The receivers use the same circuit configuration as the super buffer to drive the typically long interconnect from the chip periphery to the core. Fig. 13 shows a standard cell test chip with single-ended and differential I/O cells, a toggle flip-flop, and a

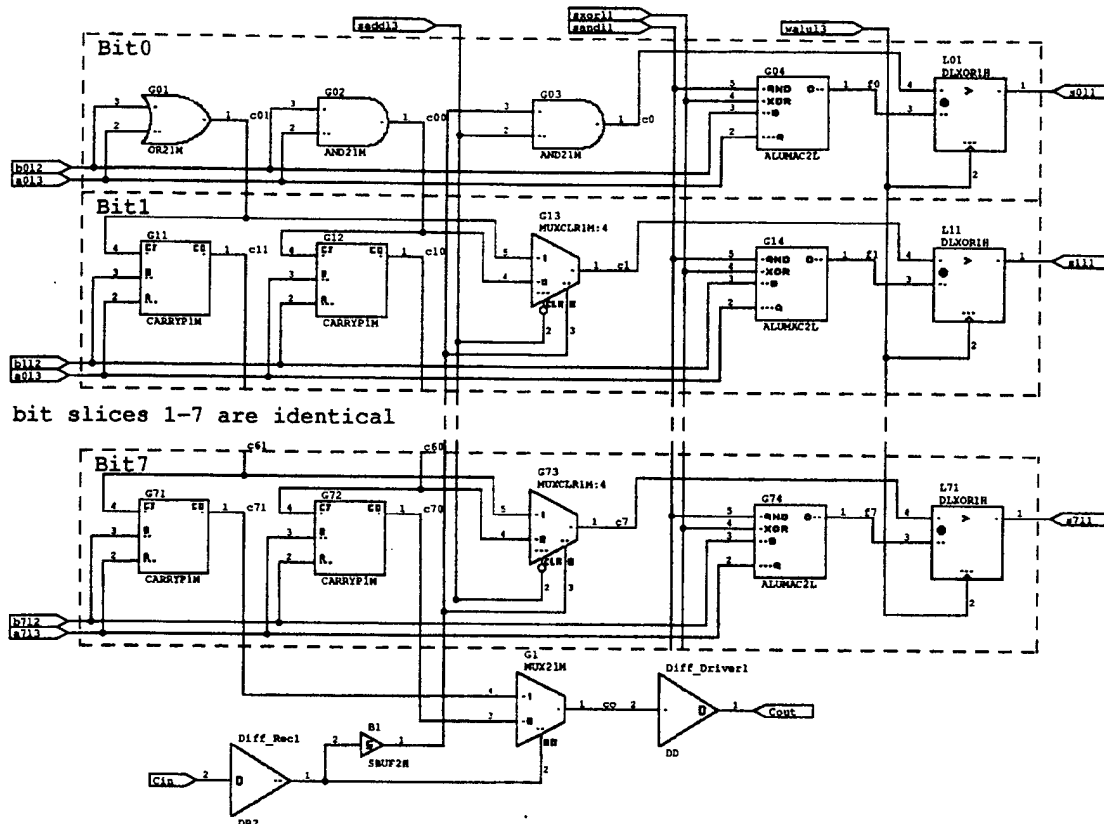


Fig. 14. 8-b carry select ALU slice.

structure to measure interconnect delays of ECL buffers, and a bias voltage generator circuit.

F. ALU Circuit

The 32-b ALU is on a critical path of FRISC since the data path had to be partitioned into four 8-b slices. The ALU has a 3-ns time slot to produce a 32-b result from the arrival of the level-2 operand. The carry select scheme is used to speed up carry propagation. The carry for each slice is calculated in two parallel carry chains, one for an assumed carry-in of one and the other for a carry-in of zero. The actual carry-in of the slice selects only the result of the appropriate carry chain. This reduces the fall-through time for the carry to a receiver, multiplexer, and driver delay if the carry chains have had time to settle. Further, the carry-in signal of the first slice must only be available on chip when the carry chains have settled. The carry select ALU can be implemented with only five current trees per bit as shown in Fig. 14.

The carry propagate gate CARRP1M and the multiplexer with clear MUXCLRIH are medium-power (10 mW) gates since they are on the critical path but drive only short interconnect. The programmable function gate ALUMAC2L generates the Boolean XOR, OR, OR AND

function of the two operands. A low-power gate (6 mW) is used since it is not on a critical path. A high-power gate (14 mW) is used for the data latch with XOR inputs DLXOR1H since it has to drive long interconnect and is on a critical path. Differential I/O drivers and receivers are used to minimize the carry fall-through time.

The ALU can perform ADD, AND, OR, and XOR functions. A subtraction is performed by inverting the carry-in and operand *B*. The output latch DLXOR1H not only latches the result but also generates the sum by performing an XOR of the carry and the XOR of the two input operands generated by the ALUMAC2L gate. Table II shows worst-case propagation delays for a 32-b add based upon SPICE simulations.

The simulation results include an average on-chip interconnect length of 600 μm between the clusters of cells that form a bit slice. The carry-in receiver and carry-out driver are placed right next to each other to avoid routing the carry-in signal all the way across the chip. The four data-path slices are mounted right next to each other on a multichip module. The off-chip interconnect between slices is at most 8 mm long. The microtransmission lines on the multichip module have a polyimide dielectric with an ϵ_r of 3.2 resulting in a low interconnect delay of 6 ps/mm. The 32-b ADD delay is the silicon delay plus

three chip-to-chip interconnect delays (3×48 ps) resulting in a worst-case delay of 2.79 ns. Assuming the clock skew can be controlled within ± 100 ps the ALU can perform a worst-case 32-b ADD within the allocated 3-ns time slot. By using carry select over a group of 3 and then 5 b the delay of the first slice could be reduced to 850 ps, resulting in a worst-case delay of only 2446 ps.

G. Standard Cell Library

The following list shows the differential standard cells used for the FRISC project. Many cells map into dual logic gates like AND and OR. Dual cells are available in the schematic library but are mapped onto the same cell during netlist expansion. Further, every input and output port of a differential cell can be inverted at no cost. Most cells are available with three different power levels ($\langle p \rangle$): low power = 6 mW, medium power = 10 mW, high power = 14 mW and with three different output levels ($\langle l \rangle$): level 1, level 2, level 3). Master/slave latches dissipate an additional 2 mW. The library also includes a 32×8 -b single-port memory cell for the register file of the processor [3], [14].

Combinational Cells

AND2($\langle p, l \rangle$)	dual-input AND gate
XOR2($\langle p, l \rangle$)	dual-input XOR gate
AND3($\langle p, l \rangle$)	three-input AND gate
XOR3($\langle p, l \rangle$)	three-input XOR/full adder
COMP($\langle p, l \rangle$)	comparator with enable
ANDOR($\langle p, l \rangle$)	AND/OR gate
ALUMAC($\langle p, l \rangle$)	programmable AND/XOR/OR gate
CARRY($\langle p, l \rangle$)	carry propagate gate

Multiplexer Cells

MUX2($\langle p, l \rangle$)	dual-input multiplexer
MUXCLR($\langle p, l \rangle$)	dual-input multiplexer with clear
MUX4($\langle p, l \rangle$)	four-input multiplexer

Buffers and Level Shifters

BUF($\langle p, l \rangle$)	buffer
SBUFH($\langle l \rangle$)	super buffer
LS($\langle l \rangle$)	level shifter

Storage Cells

SRF($\langle p \rangle \langle l \rangle$)	set-reset flip-flop
DL($\langle p \rangle \langle l \rangle$)	simple data latch
DLC($\langle p \rangle \langle l \rangle$)	data latch with synchronous clear
DLAND($\langle p, l \rangle$)	data latch with AND gate inputs
DLXOR($\langle p, l \rangle$)	data latch with XOR gate inputs
DLMUX($\langle p, l \rangle$)	data latch with MUX gate inputs
MSL($\langle p, l \rangle$)	master/slave latch
MSAND($\langle p, l \rangle$)	master/slave latch with AND gate inputs
MSMUX($\langle p, l \rangle$)	master/slave latch with MUX gate inputs

I/O Cells

SEDS	single-ended driver ECL 10K
SER	single-ended receiver ECL 10K
DD	differential driver
DR	differential receiver

Special Cells

RF32 \times 8	32 \times 8-b memory cell
SYNC	four-phase clock generator

TABLE II
WORST-CASE SILICON DELAYS FOR 32-b ADD

Chip/Circuit	Path	Delay
SLICE 1	A Op \rightarrow C _{out} _s11	1.196 ns
SLICE 2	C _{out} _s11 \rightarrow C _{out} _s12	0.451 ns
SLICE 3	C _{out} _s12 \rightarrow C _{out} _s13	0.451 ns
SLICE 4	C _{out} _s13 \rightarrow Sum_32	0.550 ns
32-b ALU	A Op \rightarrow Sum_32	2.648 ns

TABLE III
TYPICAL LOGIC DELAYS

Current Switch Delay	45 ps
Level-1 Output	40 ps
Level-2 Output	45 ps
Level-3 Output	55 ps
Fan-out Penalty per Current Switch	5 ps

A simple delay model is given to the designer which allows quick evaluation of different circuit configurations. Table III gives approximate delay figures for the current switches and the emitter followers.

The fan-out penalty for a medium-power gate is only 5 ps. However, gates like the four-input multiplexer shown in Fig. 8 can have two current switches connected to the same cell input port. Only one of the current switches can, however, be active. A detailed delay model will be described in the following section. Table IV shows typical interconnect delays.

III. MODELING OF DIFFERENTIAL CURRENT TREE LOGIC

The design of high-speed digital circuits relies heavily on accurate circuit simulation to detect problems and predict performance before fabrication. For simulation at the circuit level, SPICE provides excellent results, however, its simulation speed is prohibitively slow for large digital circuits. Digital simulators use simple digital models and event-driven timing control [15], which allows simulation of very large circuits. However, most simulators are geared towards CMOS because of its dominance in the market place. As described in [16], single-ended bipolar transistor subcircuits can be mapped into equivalent logic gates that can be simulated on a conventional digital simulator. Another modeling technique transforms the transistor-level circuits into labeled weighted graphs [17] requiring a highly specialized simulation tool.

The model presented here uses a current switch, two or more transistors connected at a common-emitter node, as a model primitive, and allows the simulation of either differential or single-ended circuits. Only the mapping of transistors with a common-emitter node to a current switch is required to generate a simulation model from a device-level (SPICE) description. The structure of the tree models is the same as the physical structure. The current switch primitive can easily be added to digital simulators that support user extensions.

TABLE IV
INTERCONNECT DELAYS

Offset	Low-Power Gate 6 mW, $I_{ef} = 400 \mu A$	Medium-Power Gate 10 mW, $I_{ef} = 800 \mu A$	High-Power Gate 14 mW, $I_{ef} = 1.2 mA$
level 1	48 ps/mm	29 ps/mm	24 ps/mm
level 2	62 ps/mm	35 ps/mm	26 ps/mm
level 3	74 ps/mm	42 ps/mm	31 ps/mm

High-speed current tree logic has several properties that needed to be modeled. The signal path and therefore the delay from an input to the output can depend upon input signals at higher levels in the tree. Thus the propagation delays from a certain level input to the output can depend upon the state of other input signals. A simple behavioral model cannot capture these delay dependencies since the delays are calculated in most simulators before the simulation starts. However, they can easily be captured by a structural model based on current switches since the actual signal path through the tree is simulated. The current switch primitive can be described with a simple behavioral model that is easy to implement on most digital simulators. The output of a current tree can be independent of a signal at a lower level. For example, if the lowest input signal of an AND current tree is undefined, the output should still be low if any of the other input signals is low. This is very important because most digital simulators set all nodes initially to the undefined state, "x." Further, the treatment of glitches is important for latches. Clock signals generated by a gate with an unsymmetrical tree have short glitches at each differential signal transition. The two signals of a differential pair are both low or both high during the glitch. Latches must be able to capture valid data if the necessary setup and hold times have been observed, even if such glitches occur on clock lines.

A. Digital Current Switch Model

Asymmetrical current trees have nonsimultaneous output signal transitions and transient glitches. The output signals of a tree can be equal during transients even though no output change should occur according to the truth table. If such glitches occur on clock lines latched data can be disturbed. The current switch model must, therefore, handle differential and nondifferential input signal conditions as shown in Fig. 15.

The simulation of differential logic on the current switch level increases the number of nodes and elements in the netlist and will slow down simulations. However, the slower simulation time must be traded off against increased accuracy and the ability to capture transients which might affect circuit performance.

Simulation efficiency could be improved by representing each differential signal pair with a single digital node. The two differential current tree outputs (q, qb) can be converted into a single-ended signal with a differential-to-single-ended converter. This converter marks nondifferential outputs of the current tree with an unknown

Extended Truth Table

Inputs			Outputs	
COM	IN	INB	Q	QB
1	?	?	1	1
X	?	?	X	X
0	1	0	0	1
0	0	1	1	0
0	1	1	0	0
0	0	0	0	0
0	X	X	0	0
0	1	T	0	1
0	0	T	1	0
0	T	T	0	0
0	X	T	0	0

X=undefined, T=threshold, ?={0,1,X,T}

COM,Q,QB represent current levels (low=on)

Signal Strength: low > high




Fig. 15. Digital current switch model.

TABLE V
PARAMETERS OF CURRENT SWITCH MODEL

C_C	10 fF
C_B	40 fF
C_E	0 fF
T_d	$40 \text{ ps} + 400 \Omega \cdot C_{\text{load}}$

logic signal. The current switch can be reduced for such a single-ended simulation of differential circuits to a four-terminal device. The single-ended modeling of differential signals reduces the number of nodes, but it requires inverter primitives for differential signal inversion. In the unlikely case that each gate output signal needs to be inverted, the total number of nodes will be larger due to the additional converters. The single-ended modeling of differential signals makes probing and saving of simulation results more efficient and allows the use of standard fault simulation and test-pattern generation software.

Negative logic is used to represent a current flowing in or out of the common-emitter node or the q and qb output nodes. Both outputs are active if the two inputs are equal and current is flowing into the common-emitter node. Latches would lose data just copied if the current switch connected to the clock signal would output no current for nondifferential inputs. However, if it outputs current on both sides no data is lost as long as the input current switch and the feedback current switch outputs agree. This will be the case as long as the data input is stable. The model will therefore correctly indicate a longer hold time and not loss of data. Sending current on both outputs for a nondifferential input condition reduces




Fig. 16. SPICE and digital XOR3T model.




Fig. 17. SPICE and digital AND3T model.


modeling pessimism in general since the tree output might not depend upon which way the current flows for a given set of input signal states.

The current switch model uses only a simple inertial delay model with capacitive load delays. If at least four signal strength values are available, the signal strength can be used to mark signal levels, which allows the detection of level violations causing saturation. The model includes capacitors to model input and output loading. It assumes differential input signals since the base capacitors are physically between base and emitter and can only be modeled as shown for differential input signals. However, most digital simulators support only capacitors connected to ground. The model parameters (Table V) depend on the operating conditions of the current switch like the switching current I_s , the voltage swing at the output, and V_{CB0} of the transistors. The capacitor C_B is 20% larger in an active current switch. This represents a

dynamic load change that is hard to simulate. However, it will be shown that a simple current switch model for all three levels can give excellent results since the dependencies are intrinsically small. In order to see the small differences the digital simulator would have to be run with a time step Δt below 5 ps. The match to a current switch at level 1 is most important since long and therefore critical signals are routed preferably on the topmost level to reduce propagation delays. The biggest simulation error is introduced by using a fixed C_B . Table V shows the parameters for the current switch with $I_s = 400 \mu A$, $V_i = 250$ mV, and $V_{CB0} = 0.85$ V.

B. Modeling of Current Trees

Figs. 16 and 17 show a comparison of SPICE simulation results with a digital simulation of a three-input xor tree and an AND/OR tree. The match for the symmetrical xor



SPICE Model: in: da,ab db,bb dc,cb out: dq,qb
DIGITAL Model: in: da,ab db,bb dc,cb out: dq,qb

Fig. 18. Glitch of three-input AND tree.

TABLE VI
SPICE AND FASTSIM RESULTS WITH INTERCONNECT DELAYS

PATH	SPICE	FASTSIM $\Delta t = 1$ ps	FASTSIM $\Delta t = 5$ ps	FASTSIM $\Delta t = 25$ ps
A Op \rightarrow C _{out}	1196 ps	1185 ps	1185 ps	1300 ps
C _{in} \rightarrow C _{out}	451 ps	448 ps	450 ps	475 ps
C _{in} \rightarrow Sum	550 ps	567 ps	570 ps	550 ps

(full adder) gate is excellent. The asymmetry of the AND tree results in nonsimultaneous output transitions that show up as rise-time degradation in the SPICE output as shown in Fig. 17. The match for the asymmetrical AND gate is clearly not as good as for the XOR gate.

Fig. 18 shows a characteristic glitch of the AND/OR current tree if the lowest level signal goes high with level-2 input high and level-1 input low. One might not expect an output transient for an AND gate with one input kept at a static low. However, the current has to propagate through two current switches after the level-3 input transition before the q output is pulled low again. SPICE shows only a signal level degradation, which can, however, lead to erroneous switching in a noisy environment. The digital model marks the glitch with nondifferential outputs allowing the detection of circuits that are sensitive to these transients. All current trees with unequal path delays from a particular current switch to the output show similar glitches. The three-input AND gate represents the worst case and should be used with caution on clock lines.

C. Modeling Accuracy Compared to SPICE

For verification of the accuracy of the standard cell models, the ALU slice shown in Fig. 14 was modeled with SPICE and FASTSIM, a digital simulator from Tektronix. Current switch and level-shifter primitives were added through its C-language interface. Table VI shows that excellent agreement (4% deviation) is possible. However, the digital simulator must be run with a sufficiently small

time step (5 ps) to avoid the accumulation of rounding errors. The delay sensitivities towards interconnect capacitance were extracted from SPICE data by a six point linear regression analysis in the range 0–500 fF.

IV. CONCLUSION

An experimental standard cell library with a typical gate delay of 90 ps for a 10-mW gate has been developed. High performance is achieved by combining advanced bipolar technology and differential current tree logic design. Interconnect delay sensitivities have been reduced by using low differential logic swings of ± 250 mV and improved ECL output drivers and buffers. Power and performance has been improved by providing different output drivers for each cell such that speed versus power can be traded off for every signal.

I/O delays can be significantly reduced by using high-speed differential drivers with low logic swing of ± 250 mV and multichip packaging. Differential I/O circuits consume further less power and are balanced, thereby avoiding delta- I noise problems.

Modeling differential logic at the current switch level gives excellent delay accuracy and allows the designer to capture transients and glitches that could cause circuit failure. Further, the modeling approach can be implemented on conventional digital simulators.

ACKNOWLEDGMENT

The authors would like to thank Tektronix for their support of the FRISC project and the fabrication of the advanced bipolar circuits.

REFERENCES

- [1] T. Yamaguchi *et al.*, "Process and device performance of a high-speed double poly-Si bipolar technology using borosenic-poly process with coupling-base implant," *IEEE Trans. Electron Devices*, vol. 35, no. 8, pp. 1247–1255, Aug. 1988.

- [2] H. J. Greub, J. F. McDonald, and T. Creedon, "Architecture of a 32bit fast reduced instruction set computer (FRISC) for implementation with advanced bipolar differential logic and wafer scale hybrid packaging," in *VLSI 87*. Amsterdam: North Holland, 1988, pp. 275-287.
- [3] H. J. Greub, "FRISC—A fast reduced instruction set computer for implementation with advanced bipolar and hybrid wafer scale technology," Ph.D. dissertation, Rensselaer Polytech. Inst., Troy, NY, Dec. 1990.
- [4] M. Franz *et al.*, "SH100E 10,000 gate ECL/15,00 gate array family with ECL/TTL I/O compatibility," in *Proc. IEEE CICC*, 1988.
- [5] M. P. Depey *et al.*, "A 10K-gate 950MHz CML demonstrator circuit made with a 1- μ m trench isolated bipolar silicon technology," *IEEE J. Solid-State Circuits*, vol. 24, no. 3, pp. 552-557, June 1989.
- [6] H. K. Park *et al.*, "High-speed polysilicon emitter base bipolar transistor," *IEEE Electron Device Lett.*, vol. EDL-7, no. 12, pp. 658-660, Dec. 1986.
- [7] V. E. Garuts, E. O. Traa, Y.-C. S. Yu, and T. Yamaguchi, "A dual 4-bit, 1.5Gs/s analog to digital converter," presented at the Bipolar Circuit and Technology Meeting, Sept. 1988.
- [8] R. L. Treadway, "DC analysis of current mode logic," *IEEE Circuits and Devices*, pp. 21-35, Mar. 1989.
- [9] W. Fang, "Accurate analytical delay expressions for ECL and CML circuits and their applications to optimizing high-speed bipolar circuits," *IEEE J. Solid-State Circuits*, vol. 25, no. 2, pp. 572-583, Apr. 1990.
- [10] H. Ichino *et al.*, "Super self-aligned technology (SST) and its applications," presented at the Bipolar Circuit and Technology Meeting, 1988.
- [11] M. Suzuki, M. Hirata, and S. Konaka, "43-ps 5.2-GHz macrocell array LSI's," *IEEE J. Solid-State Circuits*, vol. 23, no. 5, pp. 1182-1187, Oct. 1988.
- [12] S. Shimizu *et al.*, "An ECL-compatible GaAs SCFL design method," *IEEE J. Solid-State Circuits*, vol. 25, no. 2, pp. 539-545, Apr. 1990.
- [13] L. G. Heller and R. Griffin, "Cascode voltage switch logic: A differential CMOS logic family," in *ISSCC Dig. Tech. Papers*, Feb. 1984, pp. 16-17.
- [14] H. J. Greub, J. F. McDonald, and T. Creedon, "Key components of the fast reduced instruction set computer (FRISC) employing advanced bipolar differential logic and wafer scale multichip packaging," presented at the Bipolar Circuit and Technology Meeting, 1988.
- [15] E. G. Ulrich, "Exclusive simulation of activity in digital networks," *Commun. ACM*, vol. 12, no. 2, pp. 102-110, Feb. 1969.
- [16] P. Kozak, A. K. Bose, and A. Gupta, "Design aids for the simulation of bipolar gate arrays," in *Proc. 20th DAC*, June 1983, pp. 286-292.
- [17] I. N. Hajj and D. Saab, "Switch-level logic simulation of digital bipolar circuits," *IEEE Trans. Computer-Aided Design*, vol. CAD-6, no. 2, Mar. 1987.



Hans J. Greub received the diploma in electrical engineering from the Swiss Federal Institute of Technology in December 1983. He received the Master's degree in 1985 and the Ph.D. degree in 1990 from Rensselaer Polytechnic Institute in Troy, NY. He studied at Rensselaer Polytechnic Institute with the aid of a Fulbright and a Tek Labs Scholarship. During his Ph.D. studies he spent a year at Tektronix, Inc. in Beaverton, OR, working on advanced bipolar circuit design.

Since November 1990 he has been a Visiting Assistant Professor at Rensselaer Polytechnic Institute. His research interests include high-speed digital circuits, logic design, and reduced instruction set architectures. He holds three patents.



John F. McDonald (S'63-M'71) was born on January 14, 1942 in Narberth, PA. He received the B.S.E.E. degree from the Massachusetts Institute of Technology, Cambridge, in 1963, the M.Eng. degree from Yale University, New Haven, CT, in 1965, and the Ph.D. degree from Yale University in 1969.

He spent a brief period as a Member of the Technical Staff at Bell Laboratories in 1965. He held the position of lecturer at Yale University in 1969 and was appointed Assistant Professor there in 1970. In 1974 he joined Rensselaer Polytechnic Institute, Troy, NY, as an Associate Professor, where he is currently a full Professor in the Department of Electrical, Computer and Systems Engineering. He was one of the founding members of the RPI Center for Integrated Electronics in 1980. His publication record consists of roughly 150 articles, approximately a third of which are journal articles. He holds six patents, and an additional eight disclosures. His interests include VLSI and computer design, with emphasis on very high-speed electronic packaging and GaAs RISC processors.

Dr. McDonald is listed in 24 compilations of technical recognition, including *Who's Who in the World* and *American Men and Women of Science*.

Ted Creedon was with Tektronix Research Labs in Beaverton, OR. He is now President of Kestrel Technologies in Lake Oswego, OR. Kestrel manufactures R&D RF probe stations and has developed low-cost RF probes for 1-ns FRISC device and slice characterization.



Tadanori Yamaguchi received the B.S. degree in electrical engineering from Miyakonojo Institute of Technology, Japan, in 1969.

In 1969 he joined Sony Corporation, Japan, where he worked on wafer processing and device characterization of high-frequency bipolar and MOS transistors. From 1971 to 1976 he was engaged in research and development of exploratory MOS devices and process technologies. His main accomplishments were non-volatile MAOS memory devices and integrated circuits (MAOS EPROM), continuously variable threshold-voltage MOS devices, ion implantation for MOS devices, resistive-gate charge-transfer devices, and CMOS integrated circuit technology using SIPOS film. Since 1977 he has been with Tektronix, Inc., Beaverton, OR, where he developed a new approach for submicrometer-channel NMOS device technology, a 1000-V NMOS-IC technology, a 2- μ m double-layer metal CMOS-VLSI technology, latch-up-free submicrometer-channel CMOS devices using deep-trench isolation and self-aligned TiSi₂ technologies, and worked on small-geometry MOS device characterization and modeling. Currently, he is a manager of the Advanced Technology Development Department at Tektronix, where he has developed and directed high-speed self-aligned double-polysilicon bipolar and BiCMOS technologies and integrated circuits. He holds eight U.S. patents, and has published and presented over 60 technical papers including several invitational seminars.

Design of a 32 b Monolithic Microprocessor Based on GaAs HEMSFET Technology

Chien-Kuo V. Tien, Kelvin Lewis, Hans J. Greub,
Tom Tsen, and John F. McDonald

Abstract—This paper examines the design of a 32-b GaAs Fast RISC microprocessor (F-RISC/I). F-RISC/I is a single chip GaAs HEMSFET processor targeted for implementation on a multichip module (MCM) together with cache memories. The CPU architecture, circuit design, implementation, and testing are optimized for a seven-stage instruction pipeline implemented with GaAs super-buffered FET logic (SBFL). We have been able to verify novel GaAs SBFL standard cells and compare measured CPU performance with performance estimates based on circuit and device models. The prototype 32-b microprocessor has been implemented using an automated standard cell approach because of time constraints and fabricated using an experimental process by Rockwell International. The CPU chip integrates 92 340 transistors on a $7 \times 7 \text{ mm}^2$ die and dissipates 6.13 W at 180 MHz. Test results from a prototype fabrication run have demonstrated the operation of the ALU, the program counter, and the register file with delays below 6.5, and 3.4 ns, respectively. The successful modeling and verification indicate that a $0.5 \text{ }\mu\text{m}$ HEMSFET implementation of F-RISC/I could achieve a peak performance of 350 MHz. The wiring delays account for 42% of the critical path delay.

Index Terms—GaAs HEMSFET, instruction pipeline, microprocessor design, multichip module (MCM), reduced instruction set computer (RISC), super-buffered FET logic (SBFL).

I. INTRODUCTION

Recent advances in GaAs Heterojunction MESFET (HEMSFET) technology have led to gate delays below 100 ps [1] and higher integration levels, reaching VLSI complexity and, thereby, allowing the implementation of a 32 b GaAs RISC on a single chip [2]. However, integration levels are still very low compared to CMOS and do not allow the inclusion of sufficiently large caches on the chip. The cache memories must be implemented with high speed SRAM chips which need to be placed close to the CPU chip on an MCM to keep the interconnect delays low. The processor design, therefore, must consider the interactions between architecture, circuit technology, and MCM packaging. The main issues in GaAs microprocessor design are the processor versus memory speed mismatch and the limited off-chip communication bandwidth.

To overcome the difficulties of limited yield and low I/O bandwidth in GaAs, the high speed processing node, consisting of the processor and cache memory hierarchy, must be densely implemented on an MCM [3], [4]. F-RISC/I employs further a pipelined cache memory access [5] to "hide" some of the chip-to-chip delays in pipeline stages since, even on an MCM, the address and data transfer times between chips are of the same order as the processor delays.

Manuscript received March 16, 1996; revised July 29, 1996. This work was supported in part by the IBM T. J. Watson Research Center and Rockwell International and also in part by the companion F-RISC/G Research under ARPA/ARO Contract DAAL03-90-G-0817.

C.-K. V. Tien, H. J. Greub, and J. F. McDonald are with the Center for Integrated Electronics, Rensselaer Polytechnic Institute, Troy, NY 12180 USA.

K. Lewis is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.

T. Tsen is with the Microelectronics Technology Center, Rockwell International Corporation, Newbury Park, CA 91320 USA.

Publisher Item Identifier S 1063-8210(97)00741-5.

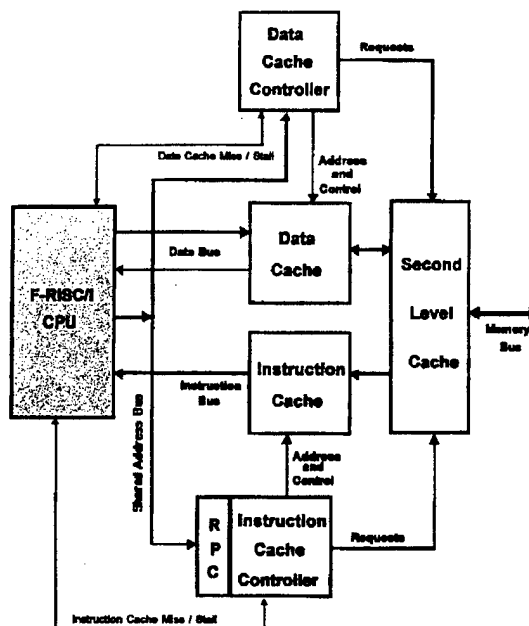


Fig. 1. FRISC/I MCM system.

The primary goals of the Fast RISC/I (F-RISC/I) project were to verify the novel GaAs SBFL standard cells, to verify that HEMSFET yields have reached adequate levels, and to correlate measured CPU performance with simulations based on circuit and device models to check the modeling capabilities of our CAD tools. F-RISC/I is a companion project to the ARPA sponsored heterojunction bipolar transistor (HBT) F-RISC project.

II. GAAS MICROPROCESSOR DESIGN

Fig. 1 shows the F-RISC/I MCM system. High memory bandwidth is achieved using separate instruction and data caches with their own data buses. This allows one 32 b instruction and one 32 b data word to be supplied by the cache memories in each cycle. A shared address bus is used to communicate with both the instruction and data cache memories to reduce CPU pin count and interconnections. This requires a remote program counter (RPC) on the instruction cache controller. Using the RPC the instruction cache can access consecutive instructions without an address transfer from the CPU. The instruction memory needs an address from the CPU only if a branch or an exception is taken. The shared address bus never causes contention in this scalar architecture since load/store and branch instructions are designed to use the address bus in the same pipeline stage.

The relative performance figure of a processor implementation is usually expressed in MIPS (millions of instructions per second), and is inversely proportional to the cycle time T_{cycle} and the average number of Cycles per Instruction (CPI). The principal parameters affecting T_{cycle} are the GaAs circuit technology and the pipeline. The CPI for a scalar RISC processor is basically one instruction cycle plus the average number of wasted cycles due to pipeline hazards, such as branch and load penalties, and stall cycles after cache misses. The CPU/system performance must be optimized by considering both the architecture (especially the pipeline scheme) and




Fig. 2. Potential instruction pipeline schemes for F-RISC/I.

technology parameters (GaAs MESFET and MCM technology) since a change in the design parameters can affect cycle time, pipeline dependencies, and/or cache miss rate in opposite ways. Since F-RISC/I is a processor without pipeline interlocks the load and branch latencies are visible at the architecture level and thus the pipeline depth is not just an implementation issue. In order to select the most advantageous pipeline for the given circuit, memory, and packaging technology, a first order performance evaluation of different pipeline schemes is necessary [5].

The evaluation starts by prototyping the critical circuits, such as the ALU and register file for the HMESFET circuit technology. Then chip placement, wiring rules, I/O pad, and layer assignments as well as propagation delay models are formulated. The MCM characteristics are: delay = 10.4 ps/mm, thermal resistance = 14°C/W, maximum power density = 8 W/cm². The cache memory design is based on a 1.5 ns 4 k × 16 b BiCMOS memory with a power dissipation of 4.5 W described in [6].

The potential instruction pipelines, shown in Fig. 2, were considered and evaluated based on cycle time and CPI. The nine-stage pipeline allocates one full cycle for address/data transfer and one CPU cycle for cache memory access to allow a larger cache size and to minimize cache miss penalties. The seven-stage pipeline can also achieve a cycle time set by the GaAs technology and also provides a pipelined cache memory access, but it only provides half a CPU cycle for address/data transfers. The cycle time of the five-stage and four-stage pipelines are limited to the same cycle time, set by the cache memory access in the Data I/O (D) stage. Since the four-stage pipeline has a lower branch latency, the five-stage pipeline can not be optimal and does not need to be considered further. Although the four-stage pipeline has a longer cycle time with address/data transfer and cache memory access performed in one cycle, it has lower branch and load latencies. To get a first order estimate of the pipeline efficiency of each pipeline scheme, compiler branch and load delay slot fill-in probabilities [7] for a six-stage RISC pipeline machine are used. Table I shows the CPI contributions of branch and load latencies for a dynamic instruction mix derived from set of typical UNIX programs [8].

A longer cycle time or more cycles for address/data transfer allows the implementation of larger caches on the MCM because more

TABLE I
CPI CONTRIBUTIONS DUE TO BRANCH AND LOAD PENALTIES

Instruction	Instruction frequency ¹	Cost 4-stage	Cost 7-stage	Cost 9-stage
ALU	60%	1	1	1
BRANCH (not taken)	9%	1	1	1
BRANCH (taken) ²	15%	1.368	2.673	3.673
LOAD ³	15%	1.1	1.5	2.3
STORE	9%	1	1	1
CPI _{intrinsic}		1.07	1.326	1.596

¹ Average dynamic instruction mix [8].

² Cost of a taken branch depends upon the branch latency slot fill-in probability of the compiler. The compiler branch fill-in probabilities are taken from [7]. (4-fill = 0, 3-fill = 36.8%, 2-fill = 20.2%, 1-fill = 16.5% and 0-fill = 26.5%).

³ Cost of a load instruction depends upon the load latency slot fill-in probability of the compiler. The compiler load latency slot fill-in probabilities are taken from [5].




Fig. 3. Chip placement on MCM for seven-stage pipeline F-RISC/I.

SRAM chips can be reached within the available address/data transfer time resulting in lower cache miss penalties. However, a longer cycle time (four-stage pipeline) results in a slower peak instruction execution rate. Allowing more cycles/stages (seven- and nine-stage pipeline) for address/data transfer increases load and branch penalties. The tradeoff between cycle time, load/branch penalties, and cache miss penalties can be evaluated by comparing the relative performance among these three pipeline candidates in a spread sheet.

Based on the address/data transfer time for each pipeline candidate, a preliminary placement of the cache memory (SRAM) chips can be determined. The chip-placement is then handled in the package design phase which includes the considerations of net routability, noise tolerance, and thermal management [9], [10]. The final package design allows the system designer to evaluate the cache miss penalties primarily based on the cache organization and the size of cache memory.

The size of the first-level cache for the seven- and nine-stage pipeline primarily depends on three factors: thermal management, net routability/topology, and allocated address/data transfer time. Close placement of a large number of chips makes heat removal more challenging and more expensive. We place the chips in a two-dimensional array with a chip pitch of 8 mm on the MCM. The junction-to-ambient resistance for each chip is estimated to be 14°C/W [9], [10] and results in a junction temperature of 64°C above ambient temperature. Since the rise time of the signals on the MCM are in the range of 200–300 ps and the MCM interconnects are 50–60 Ω transmission lines, each long net must be routed in a chaining tree (no forks) to avoid reflections. For example, the address bus originating from the cache controller chip needs to be routed as a chained net across all cache memory chips. Fig. 3 shows an example of a chip placement for the seven-stage pipeline scheme.

TABLE II
PERFORMANCE VERSUS PIPELINE DEPTH

Pipeline Depth	T_{cycle} (Relative)	$CPI_{average}$	CPI_{cache_miss}	Performance (Relative)
4-stage	1	1.07	0.188	1
7-stage	0.625	1.326	0.202	1.317
9-stage	0.625	1.596	0.140	1.159

¹ Cache miss ratios are based on SPEC92 benchmark trace data [11].

In order to evaluate the average number of stall cycles due to cache misses, assumptions about the second-level cache and the main memory and memory bus bandwidth must be made. Considering implementation cost and switching noise, the bus sizes between the first-level and the second-level as well as the second-level and the main memory are fixed at 16 and 32 bytes, respectively. The second-level cache is direct-mapped, and unified. It has a size of 1 Mbyte and a block size of 64 bytes. Main memory is assumed to be infinite and two-way interleaved. The primary data cache uses write-through to keep the cache and memory coherent. The ratio between the memory cycle time and the CPU cycle time (seven-stage) for the first-level, second-level, and main memory are 1, 4, and 16. The instruction and data cache sizes that yield optimal performance given the MCM and the BiCMOS memory characteristics are (32k, 32k) for the four-stage, (64k, 64k) for the seven-stage, and (128k, 128k) for the nine-stage pipeline. We calculated the cache miss penalties using cache miss ratios from the SPEC92 benchmark suite [11] and used the published statistics [5], [7], [8] for pipeline dependencies. We used SPEC92 benchmark data for a MIPS architecture [11] which has an instruction set similar to that of the F-RISC architecture.

Table II compares the relative performance of three pipeline schemes. Clearly the seven-stage pipeline performs better than the four-stage and nine-stage pipeline. The nine-stage pipeline has the lowest cache miss penalties, but it suffers from large branch/load penalties. The four-stage pipeline has the lowest branch/load penalties and reasonably low cache miss penalties, but its longer cycle time overshadows its higher pipeline efficiency.

III. CIRCUIT DESIGN AND PROTOTYPE PERFORMANCE

Conventional MESFET devices use Schottky barriers to provide gate isolation. The logic swing of a gate is typically between 0.6–0.7 V, limited by the turn-on voltage of the gate diode. This limited logic swing places stringent requirements on the control of the threshold voltage, power rail voltage drops, temperature effects, and fan-in effects for large GaAs circuits. The HMESFET process developed at Rockwell (Fig. 4) [1] uses a thin AlGaAs layer under the gate. The Schottky barrier at the surface has the same built-in voltage as a conventional MESFET, but the forward-bias gate current is limited by tunneling through the AlGaAs barrier. The AlGaAs barrier provides a larger turn-on voltage (1.25 V), lower leakage currents, and hence a larger logic swing. The advantages gained from HMESFET logic include higher noise margin, improved performance, and lower temperature sensitivity. Most importantly, it reduces yield losses due to random threshold voltage variation.

Direct coupled FET logic (DCFL) is popular for realization of digital circuits because of its high-speed performance and low complexity. However, DCFL has a limited fan-in and fan-out capability and a high sensitivity toward capacitive loading. Therefore, DCFL usually requires more logic levels per function than CMOS [13]. In addition the nonzero voltage low (VOL) is very sensitive to E-mode threshold voltage shifts. This is aggravated when attempting to size the devices in a gate for high drive capability.




Fig. 4. Cross section of Rockwell's HMESFET device.




Fig. 5. Schematics of DCFL and SBFL NOR-2 gates.

Super-buffered FET logic (SBFL) cascades a quasicomplementary output buffer stage after the DCFL input stage. The output buffer stage improves noise margin because of the zero-voltage low ($VOL = 0$). Although the input capacitances are approximately doubled in SBFL, the higher current-drive capability still yields lower delays than DCFL. Fig. 5 shows DCFL and SBFL NOR-2 gates. The comparison between DCFL and SBFL NOR-3 gate delays as function of wire length (with fan-out = 3) and fan-out (with wire length = 0.05 mm) are shown in Fig. 6(a) and (b). An SBFL gate compared with an DCFL gate at an equivalent power level has a least twice the drive capability. SBFL has a lower delay than DCFL for a fan-out greater than four and/or interconnect wires longer than 0.1 mm. The power supply voltages for the DCFL input stage and output buffer stage are 1.6 V (V_{dd1}) and 1.2 V (V_{dd2}). The combination of V_{dd1} and V_{dd2} is chosen to make the pull-up device (Q_2) at the output buffer stage operate with a saturation current if Q_2 is turned on. The large logic swing between the DCFL and the buffer stage provided by V_{dd1} promotes the current-drive capability even further since Q_2 is in saturation with a current quadratic in $(V_{gs} - V_{th})$. Keeping V_{dd2} below the clamping voltage of 1.25 V also reduces the static power dissipation of SBFL.

Fig. 7 shows the datapath of F-RISC/I. The potential critical paths are the PC increment in I1 stage, register file reads in DE stage, ALU execution and result feed-forward in the EX stage. The longest path starts at the outputs of the result register (RES_EX), goes through the multiplexers and the ALU, and ends at the input of RES_EX. This critical path is exercised when an ADD instruction needs the result of a previous instruction.

Level sensitive scan design (LSSD) techniques [14] are used in F-RISC/I to test each submodule at-speed. The comparison between




Fig. 6. (a) Interconnect load sensitivity of DCFL and SBFL NOR-3 gates. (b) Fanout sensitivity of DCFL and SBFL NOR-3 gates.




Fig. 7. FRISC/I datapath.

simulated results and measurements is shown in Table III. Table IV shows the delay distribution on the most critical ALU path. The 32 b ALU is implemented using a two-level carry look-ahead adder with 4 b blocks at level 1. Despite the high drive capability of SBFL the interconnect delays account for 42% of the critical path delay.

The miniaturization of FET dimensions has been and continues to be the main driving force to improve circuit speed and packing density. Hence, it is desirable to predict the system performance growth with a scaled process. Based on the same layout one can evaluate the next generation F-RISC/I performance by simulating the critical paths. Using the experimental process as a benchmark, Rockwell's baseline 0.7 and 0.5 μm HMESFET process improves the K' value by a factor of 1.36 and 1.85, respectively, while the interconnect capacitance per unit length remains the same and the wire lengths scale according to published design rules [2].

TABLE III
COMPARISONS OF CRITICAL PATH DELAYS

Critical Path	Simulated Delay	Measured Delay
PC increments in I1 stage	4.21 ns	5.0 ns
Register File Read in DE stage	3.12 ns	3.36 ns
ALU Execution in EX stage	5.17 ns	6.0 ns
LSSD latch Delay	1.42 ns	1.48 ns

The critical path simulations were performed with scaled interconnect capacitance to predict an upper bound for the performance of F-RISC/I implementations. The interconnect capacitances of the automated standard cell implementation have a scale factor of one.

TABLE IV
DISTRIBUTION OF GATE AND INTERCONNECT DELAYS ON CRITICAL ALU PATH

Gate (Power Level)	Fanout	Intrinsic Delay	Interconnect Delay	Total Delay
2-Way MUX(L)	F=1	0.170ns	0.018ns	0.188ns
LSSD	F=2	0.878ns	0.538ns	1.416ns
4-Way MUX(H)	F=2	0.228ns	0.174ns	0.402ns
3-Way MUX(H)	F=2	0.187ns	0.072ns	0.259ns
2-Way MUX(L)	F=1	0.170ns	0.046ns	0.216ns
3-Input NOR(H)	F=5	0.102ns	0.128ns	0.230ns
4-Input NOR(H)	F=1	0.104ns	0.053ns	0.157ns
4-Input NOR(H)	F=5	0.126ns	0.198ns	0.324ns
3-Input NOR(M)	F=3	0.102ns	0.093ns	0.195ns
4-Input NOR(L)	F=1	0.146ns	0.120ns	0.266ns
4-Input OR(L)	F=1	0.138ns	0.102ns	0.240ns
5-Input NOR(M)	F=4	0.173ns	0.361ns	0.534ns
4-Input NOR(L)	F=1	0.146ns	0.042ns	0.188ns
5-Input NOR(L)	F=1	0.135ns	0.047ns	0.182ns
2-Input XOR(M)	F=2	0.196ns	0.176ns	0.372ns
Sum		3.001ns	2.168ns	5.169ns
Percent of Total		58%	42%	




Fig. 8. System Performance as a function of scaled interconnect capacitances.

Custom circuit design and optimized layout can reduce interconnect capacitances, and yield a scale factor below one. Fig. 8 shows the system performance as a function of scaled interconnect capacitances. The predicted maximum operating frequency for Rockwell's baseline 0.7 and 0.5 μm HMFET processes are 260–485 and 350–660 MHz, respectively. It is notable that interconnect capacitance induced delays are very significant in determining system performance. For example, HMFET gate length scaling from 0.7 to 0.5 μm yields a 35% performance improvement for the automated standard cell implementation. A 50% reduction of interconnect capacitances in the critical paths achieved through custom layout could improve performance by 25–30%. The interconnect delay accounts for 42% of the critical path delay.




Fig. 9. Probing the 212 pin F-RISC/I test chip fabricated with an experimental 0.7 μm HMFET process from Rockwell.

IV. CONCLUSIONS

We were able to verify a novel SBFL cell library including device/circuit models and have measured critical path delays of a prototype 32 b GaAs processor implemented with an experimental HMFET process from Rockwell. The 212 pin test chip shown in Fig. 9 contains 92,340 transistors on a $7 \times 7 \text{ mm}^2$ die and dissipates 6.13 W at 180 MHz. The measured delays of critical paths could be matched within 16% by simulations with HMFET SPICE models and interconnect capacitances from a three-dimensional (3-D) capacitance extraction tool [15].

Reducing interconnect capacitances would be almost as effective for improving system performance as reducing intrinsic gate delays through device scaling. A F-RISC implementation using Rockwell's baseline 0.5 μm HMFET and additional metal layers would operate between 350–660 MHz, depending on the compactness of

the layout. In order to be competitive with state-of-the-art CMOS processors an HMESFET processor would have to be implemented with at least a 0.5 μm process using full custom layout of all critical circuits and/or yields would have to improve by a factor of 4-6 to allow at least the implementation of a dual-issue superscalar RISC.

ACKNOWLEDGMENT

The authors would like to acknowledge the assistance of Cadence, Inc., of the New Jersey office proved invaluable. Special thanks are due to Dr. C. Anderson, P. Vernes, A. Cappon, and J. Toole, whose appreciation of this work made its completion possible. Testing of F-RISC/I was made possible through the use of equipment at IBM. Finally, the authors acknowledge the contributions of R. Sherburne whose work on Berkeley RISC II provided an inspiration for this project while he was a teacher at Rensselaer Polytechnic Institute, Troy, NY, in 1985.

REFERENCES

- [1] K. R. Elliott, "Heterostructure MESFET devices for VLSI applications," in *IEEE GaAs IC Symp.*, 1991, pp. 67-70.
- [2] T. Tsen *et al.*, "0.5 μm AlGaAs/GaAs MESFET technology for digital VLSI products," in *IEEE GaAs IC Symp.*, 1993, pp. 193-196.
- [3] C. K. Tien *et al.*, "F-RISC/I: Fast reduced instruction set computer with GaAs (H)MESFET implementation," in *Proc. IEEE ICCD*, 1991, pp. 293-296.
- [4] —, "F-RISC/I: System design analysis, implementation, and performance evaluation of a 32-bit microprocessor based on GaAs HMESFET technology," Ph.D. dissertation, Rensselaer Polytech. Inst., Troy, NY, 1994.
- [5] V. M. Milutinovic *et al.*, "Pipelined design tradeoffs in a 32-bit gallium arsenide microprocessor," *IEEE Trans. Comput.*, vol. 40, pp. 1214-1224, Nov. 1991.
- [6] K. Yamaguchi *et al.*, "A 1.5 ns access time, 78- μm^2 memory-cell size, 64-kb ECL-CMOS SRAM," *IEEE J. Solid-State Circuits*, vol. 27, pp. 167-174, Feb. 1992.
- [7] V. Milutinovic *et al.*, "Architecture/compiler synergism in GaAs computer systems," *IEEE Trans. Commun.*, vol. COM-35, pp. 72-93, May 1987.
- [8] M. Namjoo and A. Agrawal, "Implementing SPARC: A high performance 32-bit RISC microprocessor," *SUN Technol.*, pp. 42-48, Winter 1988.
- [9] R. R. Tummala and E. J. Rymaszewski, *Microelectronics Packaging Handbook*. Van Nostrand Reinhold, 1989.
- [10] M. Kohara *et al.*, "High thermal conduction package technology for flip chip devices," *IEEE Trans. Comp., Hybrids, Manufact. Technol.*, vol. 6, no. 3, pp. 267-271, Sept. 1983.
- [11] J. D. Gee *et al.*, "Cache performance of the SPEC92 benchmark suite," *IEEE Micro*, vol. 13, pp. 17-27, Aug. 1993.
- [12] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*. San Mateo, CA: Morgan Kaufmann, 1990.
- [13] R. Brown *et al.*, "Gallium-arsenide process evaluation based on a RISC microprocessor example," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1030-1037, Oct. 1993.
- [14] E. Eichelberger and T. Williams, "A logic design structure for LSI testability," in *Proc. IEEE Design Automat. Conf.*, 1977, pp. 462-468.
- [15] Y. L. Le Coz *et al.*, "Application of a floating-random-walk algorithm for extracting capacitances in a realistic HBT fast-RISC RAM cell," in *Proc. 11th Int. IEEE VLSI Multilevel Interconnection Conf.*, June 1994, pp. 542-544.

Correction to "Control-Flow Versus Data-Flow-Based Scheduling: Combining Both Approaches in an Adaptive Scheduling System"

Reinaldo A. Bergamaschi, Salil Raje,
Indira Nair, and Louise Trevillyan

In the above paper,¹ the first three words of the title were missing from the table of contents on the front cover.

Manuscript received March 6, 1997.

The authors are with IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.

Publisher Item Identifier S 1063-8210(97)05143-3.

¹R. A. Bergamaschi, S. Raje, I. Nair, and L. Trevillyan, *IEEE Trans. VLSI Syst.*, vol. 5, pp. 82-100, Mar. 1997.

Short Papers

Accurate High-Speed Performance Prediction for Full Differential Current-Mode Logic: The Effect of Dielectric Anisotropy

A. Garg, Y. L. Le Coz, H. J. Greub, R. B. Iverson, R. F. Philhower,
P. M. Campbell, C. A. Maier, S. A. Steidl, M. W. Ernest, R. P. Kraft,
S. R. Carlough, J. W. Perry, T. W. Krawczyk, and J. F. McDonald

Abstract—Integrated-circuit interconnect characterization is growing in importance as devices become faster and smaller. Along with this trend, interconnect geometry is becoming more complex, consisting of an increasing number of wiring levels. Accurate numerical extraction of three-dimensional (3-D) interconnect capacitance is essential for achieving design targets in the multigigahertz digital regime. Interconnect-capacitance extraction is complicated by the presence of inhomogeneous layers with differing dielectric constant. Dielectric anisotropy as well is common in many low- κ polymeric dielectrics used in high-performance IC's. A CAD procedure using the novel floating random-walk extractor *QuickCAP* is presented. Our procedure is efficient enough to extract a substantial amount of a chip's 3-D wiring. We include as well dielectric anisotropy and inhomogeneity. The procedure is not based on effective conductor geometry or on a finite-sized conductor library but rather on the entire 3-D layout, accounting for actual local variations in conductor separations and shapes. We then apply our procedure to an experimental circuit vehicle implemented in AlGaAs/GaAs heterojunction bipolar transistor current-mode logic. This vehicle is used to validate the accuracy of our CAD procedure in predicting circuit speed. Measured and predicted test-capacitor values and ring-oscillator propagation times agreed generally to within 2–4%. To verify results on a larger digital circuit, we analyzed all interconnects in an adder carry-chain oscillator using our procedure. Predicted propagation delays were generally within 3% of measurement.

Index Terms—AlGaAs/GaAs HBT ring oscillators, current mode logic, dielectric anisotropy, floating-random-walk method, IC-interconnect modeling, 3-D capacitance extraction.

I. BACKGROUND AND INTRODUCTION

Advances in digital IC technology have produced faster and smaller devices, resulting in greater integration density and improved




Fig. 1. Three-dimensional view of multilevel interconnect in a typical HBT chip. Metal 1 (M1) and Metal 2 (M2) levels are used primarily for signal routing. Metal 3 (M3) is reserved for power supply.

performance [1]. Faster devices and signal rise times have, by necessity, placed an emphasis on interconnects when attempting to control critical net propagation delay. Issues involving distributed transmission line modeling, skin-effect loss, substrate slow-wave degradation, crosstalk coupling, and, possibly, radiative electromagnetic effects must be addressed [2].

Wire resistance in scaled interconnects aggravates the propagation-delay and bandwidth problem severely, leading to the process of "reverse scaling" or "nonscaling" of interconnect cross sections. This has resulted in chips with extremely large numbers of interconnect levels—a trend that will continue. Accurate prediction of delay in complex conductor geometry requires taking into account their true three-dimensional (3-D) structure. See, for example, Fig. 1. Proper description of 3-D structure is particularly important when wiring congestion and layout geometry vary substantially. In addition, many modern interconnect fabrication processes involve inhomogeneous and possibly anisotropic dielectrics. CAD analysis of these situations can cause substantial error when using capacitance extractors incapable of handling 3-D geometry and inhomogeneous, anisotropic dielectrics.

Three-dimensional effects have been found to be important in any style of high-performance circuit design. However, these effects are particularly important in delay prediction for heterojunction bipolar transistor (HBT) circuits as might be implemented in GaAs [3], [4], SiGe [5], InP [6], or other material systems. A desirable circuit class in these technologies is full-differential current-mode logic (CML) [7]–[9]. An example of this type of circuit is shown in Fig. 2. Differential signaling helps reduce digital switching noise since the current from the power supply is held relatively constant when the input differential signals are skew free. Smaller signal swings of satisfactory noise margin are thus possible. Also, the ability to drive purely capacitive loads is enhanced. CML circuitry dissipates a relatively small amount of dynamic power. On the negative side, differential wiring doubles the number of interconnects, complicating the routing problem [10]. Operating in differential mode can also increase interconnect capacitance by placing the effective ground plane between wires. More important, as can be seen in Fig. 3, odd-mode differential signals tend to be more sensitive to horizontal electric-field lines between conductors. Unwanted horizontal, odd-mode capacitive coupling can be, unfortunately, intensified with anisotropic interlayer dielectrics (ILD's).

Manuscript received May 6, 1998; revised September 30, 1998. This work was supported in part by the SRC Center for Advanced Interconnect Science and Technology under Contracts 448.023 and 448.024 and in part by the Defense Advanced Research Project Agency under Contracts ARPA/ARO DAAH04-93-G-0477 and DARPA/ARO DAAL03-90-G-0187. This paper was recommended by Associate Editor K. Mayaram.

A. Garg and C. A. Maier were with the Center for Integrated Electronics and Electronics Manufacturing, Rensselaer Polytechnic Institute, Troy, NY 12180 USA. They are now with Advanced Micro Devices, Sunnyvale, CA 94088 USA.

Y. L. Le Coz, P. M. Campbell, S. A. Steidl, M. W. Ernest, R. P. Kraft, S. R. Carlough, J. W. Perry, T. W. Krawczyk, and J. F. McDonald are with the Center for Integrated Electronics and Electronics Manufacturing, Rensselaer Polytechnic Institute, Troy, NY 12180 USA (e-mail: mcdonald@unix.cic.rpi.edu).

H. J. Greub was with the Center for Integrated Electronics and Electronics Manufacturing, Rensselaer Polytechnic Institute, Troy, NY 12180 USA. He is now with Intel Corporation, Portland, OR 98124 USA.

R. B. Iverson is with Random Logic Corporation, Fairfax, VA 22031 USA. R. F. Philhower was with the Center for Integrated Electronics and Electronics Manufacturing, Rensselaer Polytechnic Institute, Troy, NY 12180 USA. He is now with IBM, Yorktown Heights, NY 10598 USA.

Publisher Item Identifier S 0278-0070(99)01018-0.




Fig. 2. Circuit schematic of a CML buffer and its corresponding physical layout.




Fig. 3. Electric field lines for even- and odd-mode excitations in differential coupled pairs. Dashed line represents a virtual ground plane in the odd-mode excitation (right).

This paper first presents the basis of the floating random-walk method for estimating 3-D capacitance and shows how it can be easily modified to handle uniaxial dielectric anisotropy. The second half of this paper concerns various capacitor structures and ring-oscillator circuits that have been used to validate our CAD modeling procedure for accurate prediction of electrical switching speed.

II. CAPACITANCE ANALYSIS

We present a floating random-walk method for extracting capacitance in a 3-D conductor geometry. As we have argued, the method must produce accurate estimates in large assemblages of arbitrarily shaped conductors that constitute a substantial amount of chip wiring. We begin this section with a brief review of a newly developed floating random-walk method, on which an extractor, *QuickCAP*, is based [11]. The method can include inhomogeneous dielectric media. We follow our presentation with a proof involving a simple spatial transformation. The transformation allows us to exactly account for uniaxial dielectric anisotropy by using a single effective isotropic constant along with a mathematical scaling of vertical conductor and dielectric geometry.

A. The Random-Walk Method for Calculating Capacitance [12]

The capacitance matrix of an assembly of conductors involves solution of Laplace's equation for the electric potential ψ

$$\nabla^2 \psi = 0. \quad (1)$$

The floating random-walk method efficiently solves Laplace's equation [12]. It can be used to directly extract a capacitance matrix for general-assembly conductors within a 3-D domain. Moreover, this method requires no numerical meshing, unlike conventional finite-element and boundary-integral approaches. The absence of mesh

generation is one feature enabling the efficient analysis of large numbers of conductors [13].

The electric potential at the center of a 3-D cube can be related to the potential on its surface S , provided there are no conductors or charges lying within. This center potential

$$\psi(\xi) = \oint_S d^2 \xi' G(\xi|\xi') \psi_S(\xi') \quad (2)$$

where G is the Green's function between the cube-surface point at ξ' and the center of the cube ξ .

We next consider so-called *maximal cubes*. These cubes are defined as the largest ones surrounding a point that has no conductors within it. Obviously, the largest such cube will just touch some of the conductors where the value of electric potential is *known* (in a capacitance calculation), making it possible to evaluate part of the integral (2). The remainder of the cube surface has unknown potentials. However, the unknown potentials can be in turn treated as the center points for *second-order* maximal cubes—part of the surface of which once again just touches some conductors where potentials are established. The noncontacting points on these surfaces can be used to define third-order maximal cubes, and so forth. Fig. 4 illustrates a two-dimensional (2-D) and 3-D sequence maximal squares and cubes.

Applying Gauss' law about any particular conductor, one finds the conductor charge

$$q = \varepsilon \oint_G d^2 \xi E(\xi) \cdot \hat{n}(\xi) \quad (3)$$


where G is the enclosing surface with surface points ξ . The electric field E and outward-normal vector \hat{n} are defined on the surface as well. The electric field in (3) can be expressed in terms of maximal cubes centered on G -surface points. It has been shown that [12]

$$E(\xi) = - \oint_{S(\xi)} d^2 \xi' G_E[\xi|\xi'] \psi_{S(\xi)}(\xi'). \quad (4)$$


Here, the vector Green's function G_E relates electric field to surface potential of maximal cubes $S(\xi)$ centered on the G surface points ξ . Substituting (4) into (3), and repeatedly using (2) to represent *unknown* surface potentials of higher order maximal cubes, yields, in the infinite limit, an expression for q in terms of *known* conductor potentials. To obtain the capacitance-matrix element c_{mn} between the enclosed conductor m and any other n , we set the conductor n potential to unity. Remaining conductor potentials, including that of the enclosed conductor m , are all set to zero. Our procedure results in the following infinite series for c_{mn} , that is, charge divided by unity conductor- n potential:

$$\begin{aligned} c_{mn} = & \oint_G d^2 \xi \oint_{S(\xi)} d^2 \xi' G_E(\xi|\xi') \cdot \hat{n}(\xi) \\ & + \oint_G d^2 \xi \oint_{S(\xi)} d^2 \xi' \oint_{S(\xi, \xi')} d^2 \xi'' [G_E(\xi|\xi')] G(\xi'|\xi'') \\ & + \oint_G d^2 \xi \oint_{S(\xi)} d^2 \xi' \oint_{S(\xi, \xi')} d^2 \xi'' \oint_{S(\xi, \xi', \xi'')} d^2 \xi''' \\ & \cdot [G_E(\xi|\xi') \cdot \hat{n}(\xi)] G(\xi'|\xi'') G(\xi''|\xi''') \\ & + \dots \end{aligned} \quad (5)$$

It is understood in multiple-integral series (5) that maximal-cube surfaces S coincide with the conductor- n surface. Maximal-cube surfaces \hat{S} do not coincide with any conductor. Monte Carlo evaluation of (5) defines the floating random-walk method. Walks consist of maximal cube "hops" originating with centers on the surface G



(a)



(b)

Fig. 4. Illustration of floating random walks. (a) Example of random walks in 2-D, defined by maximal squares. Shaded regions denote terminating conductors; the centers of maximal squares are labeled in consecutive numerical order. (b) A series of nested maximal cubes producing a random walk in 3-D. The surrounding gray box represents a terminating conductor; the centers of maximal cubes are labeled in consecutive numerical order.

about conductor m and terminating, eventually, on conductor n [see Fig. 4(b)]. Other elements of the conductor capacitance matrix can be found in similar fashion.

Three-dimensional capacitance extraction using the floating random-walk algorithm is efficient in a complex rectilinear geometry. The algorithm relies on, essentially, Monte Carlo evaluation of deterministic surface integrals. It typically requires only a few random-walk hops before termination. Errors are primarily statistical in nature. The algorithm evaluates electric field only at enclosing conductor surfaces—not anywhere else. No detailed numerical meshing is required to propagate random walks. Note




Fig. 5. Cross-sectional illustration of the HBT multilevel interconnects. (In our subsequent mathematical analysis, an assembly of planar dielectric layers similar to these are analyzed, each with a possibly different anisotropic dielectric constant. An example of multiple planar-layer interfaces for this subsequent analysis are clearly seen in this figure.)

as well that deterministic inversion of a linear set of equations, as with conventional finite-element and boundary-integral methods, is avoided in computing the capacitance matrix. Moreover, the algorithm is eminently parallelizable. Simultaneous random walks can be executed on separate computational nodes in a network. Remote sections of an interconnect circuit, in fact, could be analyzed simultaneously using a tiled data-base approach [12].

B. Capacitance for Anisotropic Dielectric Media

It is generally accepted that the horizontal (parallel-to-chip-plane) dielectric constant ϵ_h of many polymers, such as polyimide, differs from the vertical (normal-to-chip-plane) constant ϵ_v [14]–[18]. We now furnish a proof resulting in a simple mathematical transformation that converts a medium with a pair of uniaxial dielectric constants (ϵ_h, ϵ_v) into a single, more convenient, isotropic medium with constant ϵ' . The transformation is mathematically exact. A similar result was previously obtained by Szentkuti for the case of a microstrip transmission line [19], [20]. We show here that the technique extends to conductors in arbitrary 3-D, layered dielectric geometry.

Fig. 5 depicts laminated dielectric layers. Each layer has given uniaxial dielectric constants ϵ_h in the x and y directions and ϵ_v in the z direction. Each layer extends infinitely in x and y . The layers, of course, possibly contain conducting electrodes (interconnect wires) for which intra- and interlayer coupling capacitance is to be found.

We define within any given layer electric potential $\psi = \psi(r)$, where $r = [x, y, z]$. Outside conductors, but inside any given layer, ψ obeys the anisotropic Laplace equation

$$\epsilon_h(\psi_{xx} + \psi_{yy}) + \epsilon_v\psi_{zz} = 0. \quad (6)$$

The x, y, z subscripts denote partial differentiation. Equation (6) must satisfy intralayer conductor Dirichlet conditions and interlayer dielectric-interface conditions. We find

$$\psi_S = f(r_S), \quad \psi(r_+) = \psi(r_-), \quad \epsilon_v\psi_z(r_+) = \hat{\epsilon}_v\psi_z(r_-). \quad (7)$$

Above, ψ_S is the electric potential at conductor surfaces within the layer of interest, r_S are coordinate vectors for points on conductor surfaces, and r_+ and r_- are coordinate vectors near any dielectric interface just within and just outside, respectively, the layer of interest. We also have $\hat{\epsilon}_v = \epsilon_v(r_+)$ and $\hat{\epsilon}_v = \epsilon_v(r_-)$.

Scaling all z coordinates in (6) and (7) according to

$$z' = z\sqrt{\epsilon_h/\epsilon_v} \quad (8)$$

produces an isotropic Laplace equation. In the primed, scaled coordinates, we can write

$$\psi'_{x'x'} + \psi'_{y'y'} + \psi'_{z'z'} = 0 \quad (9)$$

where

$$\psi'_S = f'(r'_S), \quad \psi'(r'_+) = \psi'(r'_-), \quad \varepsilon' \psi'_{z,z'}(r'_+) = \varepsilon_v \psi'_{z,z'}(r'_-), \quad (10)$$

The functions $\psi(r) = \psi'[r'(r)]$ and $\psi(r_S) = \psi'[r'_S(r_S)]$. We have also defined

$$\varepsilon' = \sqrt{\varepsilon_h \varepsilon_v} \quad (11)$$

as an equivalent, isotropic dielectric constant.

We now show that transformations (8) and (11) do not change capacitance values. By definition

$$C = \frac{Q}{V} = \frac{\int_V dx dy dz [\varepsilon_h (\psi'_{xx} + \psi'_{yy}) + \varepsilon_v \psi'_{zz}]}{\int_C \psi'_x dx + \psi'_y dy + \psi'_z dz} \quad (12)$$

The total charge Q contained on the conductor is found with Gauss' law as an integral over the enclosing volume V . The integrand of the Q integral [numerator of (12)] contributes solely at conductor surfaces. The potential difference V between any electrode pair is a line integral along C . After transformations (8) and (11), we obtain our desired result, shown in (13) at the bottom of the page.

III. FABRICATION PROCESS AND TEST STRUCTURES

Validation of our modeling procedure was achieved with full-differential CML circuits fabricated in Rockwell International's high-performance HBT AlGaAs/GaAs process. Baseline HBT devices for the process have unity current-gain frequencies f_T on the order of 50 GHz in an emitter-up configuration [3]. The minimum-geometry device has an emitter area of $1.4 \times 3 \mu\text{m}^2$. For a switching current of 2 mA, unloaded gate delays on the order of 20 ps and rise times of 30–40 ps are possible [4], [21]. The test chip was fabricated on 100-mm wafers [22]. Typical HBT base widths vary from 500 to 1000 Å. Interconnect wiring levels are situated over a 25-mil-thick semiinsulating GaAs substrate, with a ground plane plated on the wafer back side.

The process provides three layers of Au-metal interconnect with a polyimide ILD shown in Fig. 5. Additional Si_3N_4 layers are used as a lower level insulator and as a top-side moisture barrier. The Si_3N_4 layer is also used as a dielectric for power-supply bypass and special analog-circuit metal-insulator-metal (MIM) capacitors. A 50-Ω/□ NiCr thin-film layer is available to implement resistors. A 25-mil-thick semiinsulating GaAs substrate lies underneath the interconnect. The thicknesses of the metal and dielectric layers are in the 0–5- μm range.¹ Second- and third-level metals are thicker than first-level to provide low-resistance power-supply busing and global-net routing.

The polyimide used in the process is DuPont 2611, which exhibits a 25% anisotropy. Dielectric anisotropy depends on the orientation of polymer chains relative to the substrate during deposition. The process provided both inhomogeneous and anisotropic dielectric properties, making it suitable for validation test structures. Using conventional capacitance extraction methods, we predicted ring-oscillator frequencies 30–40% greater than those actually observed. The initial prediction error, prior to our QuickCAP correction, consisted of a

¹Details of the Rockwell process, such as specific ILD constants and layer thicknesses, are proprietary.




Fig. 6. Microphotograph of the fabricated HBT test chip.




Fig. 7. Schematic top view of the test capacitors, showing (a) parallel-plate, (b) finger, and (c) crossover capacitors. Note that upper and lower metal layers may be any suitable combination of M1, M2, and M3.

combination of factors: some due to polyimide anisotropy and some due to the lack of a true 3-D capacitance extractor.

Validation was based on a set of test structures. The test structures included simple capacitors and ring oscillators. Ring oscillators were designed for heightened sensitivity to inhomogeneous and anisotropic ILD. Others were designed to explore typical sensitivity to 3-D conductor-geometry variability. We analyze as well a small complete logic circuit needed in an arithmetic-logic-unit (ALU) design, that is, an adder carry chain.

Fig. 6 is a die microphotograph of the test chip. It is rectangular, measuring $8.2 \times 6.1 \text{ mm}^2$. The chip contains a number of useful passive and active structures and circuits, including resistors, capacitors, inductors, transmission lines, line-coupling structures, and ring oscillators. As noted in Section I, this work concerns two types of structures: capacitors and ring oscillators.

We have performed measurements and theoretical extractions for a variety of capacitor configurations.

- **Parallel Plate**—formed by sandwiching polyimide dielectric with any two of the available three metal layers (M1, M2, M3). Fig. 7(a) illustrates the parallel-plate geometry. This geometry is useful for establishing vertical (normal-to-chip-plane) dielectric constant. A high-value MIM capacitor is also available between bottom (M1) and middle (M2) metal layers (refer to Fig. 5). The

$$C = \frac{\int_{V'} dx' dy' (dz' / \sqrt{\varepsilon_h / \varepsilon_v}) [\varepsilon_h (\psi'_{x,x'} + \psi'_{y,y'}) + \varepsilon_v (\sqrt{\varepsilon_h / \varepsilon_v})^2 \psi'_{z,z'}]}{\int_{C'} \psi'_{x'} dx' + \psi'_{y'} dy' + (\sqrt{\varepsilon_h / \varepsilon_v}) \psi'_{z'} (dz' / \sqrt{\varepsilon_h / \varepsilon_v})} = \frac{-\int_{V'} dx' dy' dz' \varepsilon' (\psi'_{x,x'} + \psi'_{y,y'} + \psi'_{z,z'})}{\int_{C'} \psi'_{x'} dx' + \psi'_{y'} dy' + \psi'_{z'} dz'} = C' \quad (13)$$




Fig. 8. Microphotograph of a ring-oscillator structure on the HBT test chip at 40 \times magnification.

MIM-capacitor dielectric consists of a thin layer of Si_3N_4 after removing any intervening polyimide.

- *Finger*—formed within the M1 layer as two interdigitated electrodes. Fig. 7(b) illustrates the finger geometry. This geometry is useful for establishing horizontal (parallel-to-chip-plane) dielectric constant. Because of space limitations on the test chip, finger capacitors were fabricated only on M1.
- *Crossover*—formed similarly to parallel-plate capacitors; the plates, however, are made up of common parallel lines, oriented so that the two plates together produce a cross-wire array. Fig. 7(c) illustrates the cross-wire geometry. This structure is useful for studying complex 3-D fringing fields likely to be encountered in modern multilevel IC's.

We now turn our attention to the ring-oscillator circuits. We have fabricated eight-stage buffer loops implemented in differential CML. Differential wires connecting the last two stages are exchanged to produce inversion feedback necessary for oscillation. Seven stages have a fanout of one, while one stage has a fanout of two. Differential input voltage levels are 0 and -250 mV. Tree current I_E for our operating circuits is 0.8 mA at a power-supply voltage $V_{EE} = -3.2$ V.

Four of the fabricated oscillators are shown in the microphotograph of Fig. 8. Both buffers and their interconnects are designated. In all, the test chip contained 28 oscillators, consisting of devices shown in the layout of Fig. 8. Devices were connected to interconnect loading structures identical to those shown in Fig. 7.

A total of six classes of 3-D buffer-interconnect environments exist on the test chip. Within any class, interconnect (interstage) wire length was 530, 1218, 1562, or 1906 μm . A single, essentially unloaded ring oscillator with an interconnect length of 15 μm was included, bringing the count to $6 \times 4 + 1 = 25$ oscillators per chip. Sufficient variation in parasitic interconnect capacitance was therefore ensured. All interconnects for any given oscillator were of identical class and length. Fig. 9 summarizes the various interconnect environments. Note that solid-electrode or parallel-wire planes may exist above the differential interconnects S and \bar{S} .

IV. MEASUREMENT AND MODELING

A. Experimental Measurement

Our test-chip wafer was divided into a square array of 25 projected reticle patterns. Each reticle contained two 8.2×6.1 mm² test




Fig. 9. (a) Schematic of an eight-stage ring-oscillator configuration. (b) Depiction of interconnect capacitive-load structures.

chips. A Summit probe station, manufactured by Cascade Microtech, equipped with a microwave coplanar probe was used for test-capacitor measurement. Probe pads were on 150-mm pitch. Single-port s -parameter measurements on the test capacitors were performed using an HP 8510C vector network analyzer, deembedding the probe parasitics. Experimental capacitance values were derived from the extracted circuit model. Our measurement approach allowed capacitance measurement as small as 0.5 pF with an accuracy of 2%.

For ring-oscillator measurements, test-chip wafers were mounted in a Tektronix probe station and secured with a water-cooled vacuum chuck. Wafer temperature was controlled with a MELCOR thermoelectric cooler fitted between wafer and chuck. Cooler surface flatness was less than 1 mil, ensuring sufficient thermal contact with the wafer. A Tektronix 7104 oscilloscope equipped with an S4 sampling head displayed ring-oscillator waveforms. Electrical contact to the test chip was provided with a standard six-channel Cascade Microtech probe (SSPGSSGPSS) pad foot print.

Test chips were designed with the Compass tool suite. PSpice and MATLAB were used for circuit simulation and data fitting. Transistor circuit models were obtained from the process design manual. They were also independently extracted by two-port s -parameter measurement with the HP network analyzer. The small amounts of wire resistance in the Au interconnections were included in the simulations.

B. Capacitors

Measured data for a selection of test-chip parallel-plate, finger (FS), and crossover (CS) capacitors were in the 0.5–10-pF range. A planar interconnect model was generated using dielectric-thickness and metal-sheet-resistivity measurements from parallel-plate capacitors and on-chip metal resistors. The numerically extracted 3-D capacitances for finger and crossover capacitors, based on our planar model, are listed in Table I. Data were obtained with a single approximate isotropic dielectric constant of 3.2 (no anisotropic correction) based on fitting to parallel-plate results. Percent differences between measured and extracted values ranged approximately 6–20%. Anisotropy obviously is not measured with these structures.

TABLE I
COMPARISON OF MEASURED AND NUMERICALLY EXTRACTED CAPACITANCE USING AN ISOTROPIC POLYIMIDE $\epsilon_v = 3.2$. FS: FINGER, CS: CROSSOVER

Capacitor	Measured Capacitance (pF)	Isotropic Model Capacitance (pF)	Difference (%)
FS1 (M1)	0.839	0.752	11.5
FS2 (M1)	0.792	0.712	11.2
CS1 (M1/M2)	1.907	1.790	6.5
CS2 (M1/M2)	2.318	2.114	9.6
CS3 (M2/M3)	1.896	1.568	20.9

TABLE II
COMPARISON OF MEASURED AND NUMERICALLY EXTRACTED CAPACITANCE USING AN ANISOTROPIC POLYIMIDE $\epsilon_h = 4.0$ AND $\epsilon_v = 3.2$. NOTE THAT ONLY FINGER AND CROSSOVER CAPACITORS ARE SIGNIFICANTLY AFFECTED BY ANISOTROPY

Capacitor	Measured Capacitance (pF)	Anisotropic Model Capacitance (pF)	Difference (%)
FS1 (M1)	0.839	0.839	0.0
FS2 (M1)	0.792	0.784	1.0
CS1 (M1/M2)	1.907	1.901	0.3
CS2 (M1/M2)	2.318	2.299	0.8
CS3 (M2/M3)	1.896	1.863	1.7

For M1 finger capacitors, capacitance primarily depends on ϵ_h for polyimide and Si_3N_4 . However, secondary fringing fields into polyimide and underlayer dielectric give capacitance contributions that depend on both ϵ_h and ϵ_v , the vertical (out-of-plane) dielectric constant. Crossover capacitance is influenced by both ϵ_h and ϵ_v for polyimide and Si_3N_4 . To obtain good agreement with extracted values for finger and M1/M2 crossover capacitors, we established a 25% polyimide anisotropy $\epsilon_h = 4.0$ and $\epsilon_v = 3.2$ [14]–[18].² Table II shows corrected Table I finger and crossover capacitance data. Correction accounts for 3-D effects using *QuickCAP* and for polyimide anisotropy using the transformations of Section II-B. Percentage differences were less than 2%. Some of these structures are used to identify the anisotropy by adjusting it for fit, while others involve prediction and measurement to confirm it.

C. Ring Oscillators

Table III is a listing of all the types of eight-stage ring oscillators in our study.³ It includes their load configuration, measured and simulated periods of oscillation, and percentage difference between theory and experiment. Circuit simulations for oscillation period relied on experimentally verified HBT device models. Interconnect capacitance was included during circuit simulation, with our previously determined $\epsilon_h = 4.0$ and $\epsilon_v = 3.2$ polyimide-ILD values. Observe that our predicted oscillation periods are within 4% of measurement. Variations in prediction accuracy with structure type are all within this same 4% range.⁴

Fig. 10 is a plot of oscillation period P versus interconnect-load capacitance C per stage. The plot displays measured and simulated data points for the 25 ring-oscillator loads on the test chip. The oscillation period can be written $P = 2Nd$ for N number of stages and d delay per stage. The plot follows the linear relation $P = kC + P_{\text{int}}$, where k is a constant slope, C is load capacitance per stage, and the intercept P_{int} is the total unloaded delay $2Nd$ when $C = 0$. Because of constant-current charging and discharging

² Throughout this paper we report the dielectric constants relative to ϵ_0 .

³ Even-stage ring oscillators can be constructed by exchanging appropriate differential signal lines.

⁴ The ring oscillators exhibited minimal temperature sensitivity over a range $\pm 30^\circ\text{C}$ about room temperature ($\sim 25^\circ\text{C}$).

TABLE III
MEASURED AND SIMULATED RING-OSCILLATOR PERIODS. POLYIMIDE ANISOTROPY WAS $\epsilon_h = 4.0$ AND $\epsilon_v = 3.2$

Oscillator	Type	Measured Oscillation Period (ps)	Anisotropic Model Period (ps)	Difference (%)
1	Unloaded	549	556	-1.35
2	Finger	812	811	0.12
3	Finger	1180	1158	1.88
4	Finger	1368	1335	2.41
5	Finger	1547	1524	1.50
6	Finger	778	785	-0.86
7	Finger	1086	1096	-0.98
8	Finger	1245	1249	-0.35
9	Finger	1390	1417	-1.90
10	Overlap	836	825	1.37
11	Overlap	1204	1184	1.63
12	Overlap	1430	1376	3.81
13	Overlap	1619	1577	2.63
14	Overlap	895	898	-0.39
15	Overlap	1385	1381	0.33
16	Overlap	1640	1620	1.22
17	Overlap	1904	1883	1.12
18	Cross-over	810	828	-2.12
19	Cross-over	1203	1190	1.15
20	Cross-over	1413	1389	1.70
21	Cross-over	1609	1572	2.29
22	Cross-over	875	881	-0.66
23	Cross-over	1336	1302	2.57
24	Cross-over	1563	1527	2.32
25	Cross-over	1791	1754	2.04




Fig. 10. Measured and simulated oscillation period versus load capacitance based on the uniaxial anisotropic dielectric model ($\epsilon_h = 4.0$, $\epsilon_v = 3.2$) for the eight-stage CML ring oscillators.

of CML interconnect capacitance, the constant $k = \Delta V / I_L$, where ΔV is the differential voltage swing (250 mV) and I_L is the tree current (0.8 mA). P_{int} solely depends on intrinsic device switching speed ($C = 0$).

D. A Practical Application: ALU Carry Chain

The polyimide dielectric model previously developed in analyzing the test chip was applied to a complex, 8-bit, ALU carry-select chain [23], [24]. It was fabricated with the same HBT reticle and process. A logic schematic of this circuit is drawn in Fig. 11. The circuit is implemented in differential CML. The chain can be set into oscillation along either short or long paths for delay measurement. The main characteristics of the carry-chain circuit are summarized in Table IV.

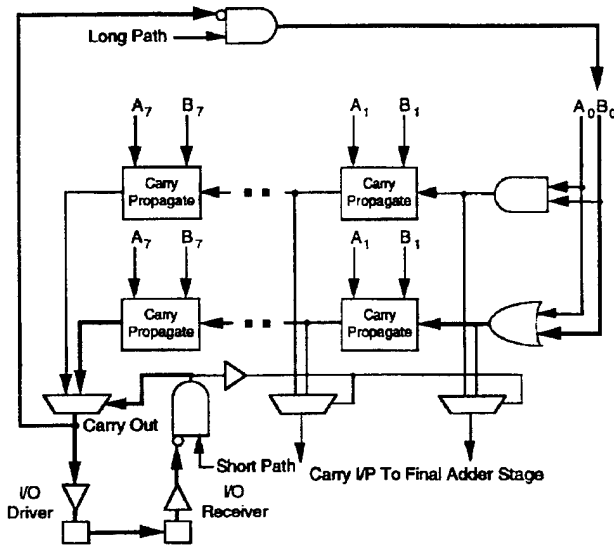


Fig. 11. Schematic of short and long paths in the self-oscillating HBT ALU carry chain. The oscillation paths are shown in bold.

TABLE IV
HBT ALU CARRY-CHAIN CIRCUIT SPECIFICATIONS

Parameter	Value
Size	1.3mm × 5.6mm
No. of Devices	658
Total No. of Nets	2037
Long-path Critical Nets	56
Short-path Critical Nets	44



Fig. 12. Layout of the ALU carry chain.

TABLE V
MEASURED AND SIMULATED ALU CARRY-CHAIN SHORT- AND LONG-PATH DELAYS. POLYIMIDE ANISOTROPY WAS $\epsilon_h = 4.0$ AND $\epsilon_v = 3.2$

Carry Path	Measured Delay (ps)	Isotropic Model (ps)	Difference (%)	Anisotropic Model (ps)	Difference
Short	520	476	8.5	505	2.9
Long	1020	936	8.2	1025	-0.5

The parasitic capacitance of 56 long-path critical nets and 44 short-path critical nets within the ALU carry chain were extracted with QuickCAP. Fig. 12 shows the layout of the analysis domain and of the extracted nets. A comparison with experimentally measured critical-path delay data is provided in Table V. The short path contains one net that is relatively long (1918 μm). It was modeled as a four-element resistance-capacitance ladder. Circuit simulations accounted for the parasitic capacitance of all the long- and short-path critical nets.

V. CONCLUSION

An AlGaAs/GaAs HBT test chip was fabricated using Rockwell International's 50-GHz baseline process. The test chip was designed

to evaluate 3-D interconnect-capacitance effects in high-speed digital circuits. This process uses polyimide ILD. The chip contained capacitor structures and ring oscillators, which were implemented in full-differential CML. A uniaxial polyimide ILD anisotropy of 25% ($\epsilon_h = 4.0$, $\epsilon_v = 3.2$) was required to fit experiment with theory (QuickCAP) by adjustment of dielectric constants for some test structures while others were evaluated to confirm the predictions. Measured test-capacitor values and ring-oscillator periods were, generally, within several percent of CAD-tool prediction. Our 25% anisotropy model was independently applied to a complex microprocessor ALU carry-chain circuit. Measured and simulated self-oscillation periods of the carry chain were within 3%.

REFERENCES

- [1] J. D. Meindl, "Gigascale integration: Is sky the limit?" *IEEE Circuits Devices Mag.*, vol. 12, pp. 19–24, Nov. 1996.
- [2] D. C. Edelstein, M. Schuermann, and L. M. Geppert, "Picosecond propagation on multilevel copper-polyimide back end of the line Interconnections," in *Proc. 10th Int. VLSI Multilevel Interconnection Conf. (VMIC-93)*, Santa Clara, CA, June 1993, pp. 511–513.
- [3] P. M. Asbeck, M.-C. F. Chang, J. A. Higgins, N. H. Sheng, G. J. Sullivan, and K.-C. Wang, "GaAlAs/GaAs heterojunction bipolar transistors: Issues and prospects for application," *IEEE Trans. Electron Devices*, vol. 36, pp. 2032–2042, Oct. 1989.
- [4] K.-C. Wang, P. M. Asbeck, M.-C. F. Chang, R. B. Nubling, R. L. Pierson, N.-H. Sheng, G. J. Sullivan, J. Yu, A. Chen, D. Clement, T. T. Tsen, H. F. Basit, J. D. George, and R. Young, "A 15-GHz gate array implemented with AlGaAs/GaAs heterojunction bipolar transistors," *IEEE J. Solid-State Circuits*, vol. 26, pp. 1669–1672, Nov. 1991.
- [5] B. S. Meyerson, "UHV/CVD growth of Si and SiGe alloys: Chemistry, physics, and device applications," *Proc. IEEE*, vol. 80, pp. 1592–1608, Oct. 1992.
- [6] B. Jalali and S. J. Pearton, Eds., *InP HBT's: Growth, Processing and Applications*. Norwood, MA: Artech House, 1995.
- [7] D. R. Breuer, "Applications of low-level differential logic," in *Dig. Tech. Papers Int. Solid-State Circuits Conf.*, 1975, pp. 126–127.
- [8] H. J. Greub, J. F. McDonald, T. Creedon, and T. Yamaguchi, "High-performance standard cell library and modeling technique for differential advanced bipolar current tree logic," *IEEE J. Solid-State Circuits*, vol. 26, pp. 749–762, May 1991.
- [9] E. B. Eichelberger and S. E. Bello, "Differential current switch—High performance at low power," *IBM J. Res. Dev.*, vol. 35, no. 3, pp. 313–320, May 1991.
- [10] J. Loy, A. Garg, M. Krishnamoorthy, and J. F. McDonald, "Differential routing of MCMs—CIF: The ideal bifurcation medium," in *Proc. IEEE Int. Conf. Comp. Des.*, Cambridge, MA, Oct. 1994, pp. 599–603.
- [11] "QuickCAP," Random Logic Corp., Fairfax, VA.
- [12] Y. L. Le Coz and R. B. Iverson, "A stochastic algorithm for high-speed capacitance extraction in integrated circuits," *Solid-State Electron.*, vol. 35, no. 7, pp. 1005–1012, 1992.
- [13] Y. L. Le Coz, H. J. Greub, and R. B. Iverson, "Performance of random walk capacitance extractors for IC interconnects: A numerical study," *Solid-State Electron.*, vol. 42, no. 4, pp. 581–588, 1998.
- [14] S. Herminghaus, D. Boese, D. Y. Yoon, and B. A. Smith, "Large anisotropy in optical properties of thin polyimide films of poly (p-phenylene biphenyltetracarboximide)," *Appl. Phys. Lett.*, vol. 59, no. 9, pp. 1043–1045, Aug. 1991.
- [15] T.-M. Lu, S. P. Murarka, T. S. Kuan, and C. H. Ting, "Low dielectric constant materials—Synthesis and applications in microelectronics," *Mater. Res. Soc.*, vol. 381, pp. 3–9 and pp. 135–139, 1995.
- [16] M. Pecht and X. Wu, "Characterization of polyimides used in high density interconnects," *IEEE Trans. Comp. Hybrids, Manufact. Technol. B*, vol. 17, pp. 632–639, Nov. 1994.
- [17] I. M. Ikeda, "A mechanical effect of orientation," *J. Polymer Sci.*, vol. 4, pp. 353–359, 1966.
- [18] R. A. Shick, B. L. Goodall, L. H. McIntosh, S. Jayaraman, P. A. Kohl, S. A. Bidstrup-Allen, and N. R. Grove, "New olefinic interlevel dielectric materials for multi-chip modules," in *Proc. 1996 IEEE MCM Conf.*, Santa Cruz, CA, Feb. 1996, pp. 182–187.
- [19] B. T. Szentkuti, "Simple analysis of anisotropic microstrip lines by a transform method," *Electron. Lett.*, vol. 12, no. 25, pp. 672–673, Dec. 1976.

- [20] L. N. Dworsky, *Modern Transmission Line Theory and Applications*. New York: Wiley, 1979.
- [21] H. Kroemer, "Theory of a wide-gap emitter for transistors," *Proc. IRE*, pp. 1535–1537, Nov. 1957.
- [22] R.-T. Huang, D. Nelson, S. Mony, R. Tang, R. Pierson, J. Penny, and R. Sahai, "Manufacturing AlGaAs/GaAs HBT's on 100 mm wafers," in *Tech. Dig. 1993 GaAs IC Symp.*, Oct. 1993, pp. 345–348.
- [23] Y. L. Le Coz, H. J. Greub, A. Garg, P. Campbell, S. A. Steidl, J. F. McDonald, and R. B. Iverson, "Interconnect-limited oscillations within a fast-RISC ALU implemented with a 50-GHz HBT process using anisotropic polyimide ILD: Simulation and measurement," in *Proc. 14th Int. VLSI Multilevel Interconnection Conf. (VMIC-97)*, Santa Clara, CA, June 1997, pp. 201–203.
- [24] Y. L. Le Coz, H. J. Greub, A. Garg, J. F. McDonald, and R. B. Iverson, "A floating random-walk algorithm for extracting IC-interconnect capacitances: Simulation and experimental verification of critical-path delay within a fast-RISC ALU," in *Proc. 13th Int. VLSI Multilevel Interconnection Conf. (VMIC-96)*, Santa Clara, CA, June 1996, pp. 230–232.

On the Design of Optimal Counter-Based Schemes for Test Set Embedding

Dimitri Kagaris and Spyros Tragoudas

Abstract—Counter-based mechanisms have been proposed for use in built-in test set embedding. A single counter or multiple counters may be used with one or multiple seeds. In addition, counters may be combined with ROM's. Each alternative design scenario introduces a difficult combinatorial optimization problem: minimization of the time required to reproduce the test patterns by an appropriate synthesis of the built-in test pattern generator. This paper presents fast synthesis techniques that result in almost optimal designs. For any given circuit, they efficiently determine whether counter-based schemes are applicable as built-in generators for a given circuit. The proposed techniques have been implemented and tested on the ISCAS'85 benchmarks. Comparative studies with a weighted random linear feedback shift register scheme show that counter-based designs may offer good hardware/time solutions.

Index Terms—Algorithms, automatic testing, delay effects, logic circuit testing.

I. INTRODUCTION

The process of built-in test pattern generation (TPG) can be separated (implicitly or explicitly) into two tasks: generation of patterns for the easy-to-detect faults and generation of patterns for the hard-to-detect faults. The first task can be easily handled with a pseudorandom pattern generator like a linear feedback shift register (LFSR). The second task is more difficult and requires some form of a *deterministic* test pattern generator. In deterministic test pattern generation, the generating mechanism has to take into account in some way each one of the specific patterns (or groups of patterns) that target the hard-to-detect faults. Below we give a brief classification of the deterministic test pattern generation methods (assuming combinational circuits and stuck-at faults with no sequential behavior).

A. A Classification of Deterministic TPG Schemes

There is a great variety of schemes that have been proposed for deterministic TPG. These schemes can be classified under different criteria, such as:

- i) *Weighting Logic—Mapping Logic*: A pseudorandom generator (typically, LFSR) is used as a basic subcomponent. The patterns generated by this generator are then transformed into the target deterministic patterns. The transformation can be done by "weighting" the bit probabilities of the pseudorandom source, or by explicitly "mapping" a subset of the pseudorandom patterns to the target deterministic patterns. Examples in the first category are [7], [20], [24], and [27] and in the second [4], [8], [28], and [30], among others.
- ii) *Test Length Bound—Fault Coverage Bound*: Some schemes give priority to not exceed a prescribed bound on the test

Manuscript received October 28, 1997; revised September 23, 1998. This research was supported by NSF Grant CCR 9815229. This paper was recommended by Associate Editor T. Cheng.

D. Kagaris is with the Electrical Engineering Department, Southern Illinois University, Carbondale, IL 62901 USA.

S. Tragoudas is with the Electrical and Computer Engineering Department, The University of Arizona, Tucson, AZ 85721 USA (e-mail: spyros@ece.arizona.edu).

Publisher Item Identifier S 0278-0070(99)01017-9.

A Very Wide Bandwidth Digital VCO Using Quadrature Frequency Multiplication and Division Implemented in AlGaAs/GaAs HBT's

Peter M. Campbell, Hans J. Greub, Atul Garg, Samuel A. Steidl, Steven Carlough, Matthew Ernest, Robert Philhower, Cliff Maier, Russell P. Kraft, and John F. McDonald

Abstract— A digital voltage-controlled oscillator (VCO) is described which uses frequency multiplication and division to achieve very wide bandwidth. The VCO uses current-mode logic and does not require reactive elements such as inductors, capacitors or varactors. A novel, fully symmetric exclusive-OR (XOR) circuit was developed which uses product pairs and emitter-coupled logic. To achieve the highest performance possible, the critical path is symmetric and special physical design techniques were developed to promote matched-capacitance. The maximum measured frequency was 13.66 GHz. The chip occupies $1.9 \text{ mm} \times 1.6 \text{ mm}$ and dissipates 2.45 W at a supply voltage of -6.0 V . With a measured frequency range from 1.25 to 13.66 GHz, this circuit has the widest bandwidth reported in the literature for any VCO, digital or analog.

Index Terms— Current-mode logic, exclusive-OR gate, heterostructure bipolar transistors, matched-capacitance layout, phase-locked loop, quadrature frequency multiplication, ring-oscillator, variable-delay element, voltage-controlled oscillators.

I. INTRODUCTION

OVER the past decade there has been an explosion in high-speed communications and electronics, most notably the advent of fiber-optic and wireless communications and the ever-increasing clock speeds of microprocessors. As the trend continues, high-speed clock synchronization and generation will become increasingly critical. Furthermore, advances in device integration leads to larger chips that require the distribution of accurate clock signals. Because the distribution tree may have radically different lengths and parasitics, synchronization of the clock signals at the receivers can be difficult and has been the subject of much interest [1], [2], [9].

Manuscript received February 3, 1997; revised October 2, 1997. This work was supported in part by the Advanced Research Projects Agency under Contracts DAAL03-90G-0187 and DAAH04-93G-0477 and AASERT Award DAAL03-92G-0307.

P. M. Campbell was with the Department of Electrical and Computer Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180 USA. He is now with IBM, Hopewell Junction, NY 12533 USA.

H. J. Greub was with the Department of Electrical and Computer Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180 USA. He is now with Intel, Portland, OR 98124 USA.

A. Garg, S. A. Steidl, S. Carlough, M. Ernest, R. P. Kraft, and J. F. McDonald are with the Department of Electrical and Computer Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180 USA.

R. Philhower was with the Department of Electrical and Computer Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180 USA. He is now with IBM, Yorktown Heights, NY 10598 USA.

C. Maier was with the Department of Electrical and Computer Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180 USA. He is now with Sun Microsystems, Palo Alto, CA 94303 USA.

Publisher Item Identifier S 1063-8210(98)02203-3.

To date, complementary metal-oxide-semiconductor (CMOS) technology has been the dominant throughout the industry, but it has not displaced other technologies for extremely high-speed designs. The use of non-CMOS and even nonsilicon technologies is often in pursuit of higher performance, either due to material properties (such as the improved electron mobility of gallium arsenide) or device technology (such as heterostructures). This improvement in performance typically has a significant cost in terms of price, reliability, or device yield.

The fast RISC (F-RISC) project at Rensselaer Polytechnic Institute, Troy, NY, was founded to investigate whether low-yield, high-speed processes could be used to create a computer with significantly higher performance. The F-RISC/G processor uses small chips on a multichip module (MCM) [3]. With a cycle time of 1 ns and a clock frequency of 2 GHz, the distribution of clock signals to all chips on the MCM is one of the most critical aspects of the design. The voltage-controlled oscillator (VCO) was developed in part to investigate the upper limits of the fabrication process selected for F-RISC/G. It also has applications in wideband phase-locked loops, communications, signal synchronization and clock generation/deskew.

II. DEVICE AND MATERIAL CHARACTERISTICS

The VCO was fabricated in the Rockwell 50 GHz baseline GaAs/AlGaAs process. Gallium arsenide (GaAs) has significantly higher electron mobility when compared to Silicon, reducing the transistor base-transit time and increasing the device performance. The use of GaAs does come with significant drawbacks such as low device integration levels and material fragility.

Heterostructure bipolar transistors (HBT's) offer significantly higher performance than homojunction bipolar transistors (BJT's). This improvement in speed is due primarily to a heterojunction at the base-emitter interface that reduces the back-injection of carriers from the base to the emitter and improves the device gain significantly. The base doping is often increased to reduce resistance at the expense of the gain and improves the performance of the device. Emitter doping can also be dropped in order to reduce the base-emitter junction capacitance C_{je} .

III. INTERCONNECT CHARACTERISTICS AND PERFORMANCE

Capacitive coupling for GaAs and other semi-insulating substrates presents a different situation than in silicon. Because the substrate is semi-insulating, the groundplane is relatively far away, reducing the capacitance to ground but increasing




Fig. 1. Voltage-controlled oscillator (VCO) architecture.




Fig. 2. VCO core signal generator and frequency multiplier.

the coupling between adjacent nodes. The anisotropic properties (larger dielectric constant in the horizontal direction) of some low K dielectrics like polyimide can also increase the wire-wire coupling, further increasing the design complexity. To avoid the large computational and memory requirements of conventional field solvers, we have used QuickCap [4] that employs random walks to extract parasitic capacitance and has significantly lower resource requirements.

IV. CIRCUIT DESIGN

The VCO (Fig. 1) is based upon a ring-oscillator composed of variable-delay elements from which multipliers generate signals at two and four times the core frequency [5]. A divider chain has been added which provides divisors of two, four, and eight. An ordinary ring-oscillator composed of 16 inverter stages is included on the chip as a process monitor.

Multiple signal paths have been provided through the circuit using multiplexers to allow the testing of partially functional chips. Only the main probe site and the high-speed output site are required for testing while the divider and external clock probe sites assume a default value when not in use. Current-mode logic (CML) is used throughout the VCO to reduce switching noise and jitter.

A. Core Oscillator and Frequency Multiplier

The core oscillator (Fig. 2) is composed of four voltage-controlled variable-delay elements that are connected in a circular fashion with one inversion along the path. The frequency is doubled and quadrupled using a novel exclusive-OR gate that is perfectly symmetric in order to reduce phase error.




Fig. 3. Voltage-controlled variable-delay element.




Fig. 4. Delay element signal path characteristics. (a) Slow and fast path device currents. (b) Signal delay for slow, combined (slow + fast), and total paths.

Because the signal has to propagate through the circuit twice to complete one oscillation, the delay of each stage is 45° . Signals separated by two stages are 90° out-of-phase and are said to be in quadrature. With the use of a multiplier, these two signals may be combined to generate a new signal at twice the original frequency. The four separate taps provide two sets of quadrature outputs that may be used to double the core frequency. Because the shift between successive stages is 45° , the phase shift between the doubled-frequency signals are also in quadrature and may be combined to produce a signal at four times the core frequency.

B. Variable-Delay Element

The core frequency is adjustable due to the use of variable-delay elements [6] (Fig. 3). Switching the current between slow and fast paths within the element varies the delay. Fig. 4(a) shows the simulated shift in current between the slow and fast paths while Fig. 4(b) shows the delay through the slow path, the combined paths (slow + fast), and the total delay including the output emitter-followers.

C. Frequency Divider

Frequency division is accomplished using sequential toggle flip-flops, each providing a successive division factor of two. Since the frequency is cut in half after every stage, only the




Fig. 5. Novel, fully symmetrical exclusive-OR circuit.




Fig. 6. Comparison of simulated waveforms from balanced Gilbert multiplier and novel XOR.

first stage requires careful design. A buffer is used between the first and second stages to ensure the quality of the signal, after which no additional buffers are required.

D. Frequency Multiplier/Novel XOR Circuit

Typically, signal multiplication is performed using analog circuits such as the Gilbert multiplier [7], which is capable of four-quadrant operation. However, in order to generate a high-quality signal at twice the input frequency, the delay for both input signals must be the same, a characteristic that the Gilbert multiplier does not possess. To compensate, Schmidt [8] combined two Gilbert cells with inverted signal connections to cancel out the input phase shift.

Our circuit does not use analog signal multiplication. Instead, the product pairs for the exclusive-OR logic function (a_0b_0, a_0b_1, a_1b_0 , and a_1b_1) are generated and combined to realize the function (Fig. 5). Although it requires the same amount of devices as the dual-multiplier approach, SPICE simulations have indicated that our circuit has higher rise times at lower frequencies (Fig. 6). Due to the perfect symmetry, the circuit has low phase error and is used as a phase detector in a 2 GHz clock deskew circuit that reduces skew to less than 5 ps [9].

V. PHYSICAL DESIGN AND LAYOUT

To produce robust multi-GHz circuits, the physical layout must be considered to be nearly as important as the circuit design itself. Mismatched parasitic loading can have disastrous effects upon the circuit performance; consequently the VCO physical layout was handcrafted. To compensate, special




Fig. 7. Closely balanced layout of VCO core, 2x and 4x multipliers with both inputs to one multiplier imbalanced.




Fig. 8. Closely balanced layout of VCO core, 2x and 4x multipliers with one input to both multipliers imbalanced.

techniques were developed to produce layouts with closely matched capacitance in order to reduce skew.

Because symmetry is essential in order to reduce phase error, the four delay elements are arranged in a square to equalize the interconnect parasitics. Two alternative arrangements for routing the signals between the delay elements and the multipliers were developed in order to investigate the effects of phase shifts upon the output signals (Figs. 7 and 8).

Assuming a sinusoidal signal generated by the core oscillator, the input to the multipliers will have a phase shift depending upon the configuration used. For the layout in Fig. 7, one set of multiplier inputs has a phase shift δ , hence the 2x and 4x multiplier outputs are (arbitrarily assigning the phase shift δ to the first set of inputs)

$$\text{OUT1} = \sin(\theta + \delta) * \sin(\theta + \pi/2 + \delta) \quad (1)$$

$$\text{OUT2} = \sin(\theta + \pi/4) * \sin(\theta + \pi/4 + \pi/2) \quad (2)$$

$$\text{OUT3} = (1/8) * [\sin(4\theta + 2\delta) + \sin(2\delta)] \quad (3)$$

For the layout in Fig. 8, the multiplier outputs are

$$\text{OUT1} = \sin(\theta) * \sin(\theta + \pi/2 + \delta) \quad (4)$$

$$\text{OUT2} = \sin(\theta + \pi/4) * \sin(\theta + \pi/4 + \pi/2 + \delta) \quad (5)$$

TABLE I
VCO STATISTIC

Chip Area	1.9 mm X 1.6 mm	Supply	$V_{EE} = -6.0$
Emitter Area	1.4 μm X 3.0 μm	Control V	± 1.0 V
Freq. Range	2.04 - 13.66 GHz	HBTs	412
Wafer Size	4-inch	Power	2.45 W




Fig. 9. Microphotograph of fabricated VCO chip.

$$\begin{aligned} \text{OUT3} = & (1/4) * [(1/2) * \sin(4\theta + 2\delta) - \sin(\delta) \\ & * (\sin(2\theta + \delta) + \cos(2\theta + \delta)) \\ & + \sin^2(2\delta)]. \end{aligned} \quad (6)$$

In (8), the phase shift is transferred to the output signal and a DC offset is generated while (11) has a subharmonic component at $2\times$. Consequently a phase shift on both inputs to one multiplier is preferable to a phase shift on one input of both multipliers, and the arrangement in Fig. 7 was selected for the VCO layout.

VI. RESULTS

A microphotograph of the fabricated chip is in Fig. 9, and the measured tuning range for the core oscillator is shown in Fig. 10. The minimum (undivided) frequency was 1.25 GHz which should result in a divided output frequency of 0.156 GHz (the divider was not tested due to a limited number of probes). The maximum measured frequency was 13.66 GHz and is shown in Fig. 11. Simulation results have indicated that the speed should be 13.9 GHz for a control voltage of 0.6 V, an error of 1.76%.

VII. SUMMARY

A voltage-controlled oscillator (VCO) with a very wide measured bandwidth of 1.25 GHz (undivided) to 13.66 GHz has been described. The circuit has the widest bandwidth reported in the literature for any VCO, digital or analog. A novel fully symmetric exclusive-OR circuit was presented and discussed along with the variable-delay frequency generation mechanism. The VCO and its subcircuits have application




Fig. 10. Measured VCO core tuning range (nonmultiplied).




Fig. 11. Oscilloscope photograph of 13.66 GHz VCO output.

in high-speed communications, phase-locked loops, signal synchronization and clock generation/distribution.

ACKNOWLEDGMENT

The authors would like to thank Y. L. LeCoz of Rensselaer Polytechnic Institute, Troy, NY, and R. B. Iverson of Random Logic Corporation for their technical assistance and support. They also thank J. Toole, B. Parker, and B. Lucas at DARPA for their encouragement and support.

REFERENCES

- [1] J. Alvarez *et al.*, "Wide-bandwidth low-voltage PLL for PowerPC microprocessors," *IEEE J. Solid-State Circuits*, vol. 30, pp. 383-391, Apr. 1995.
- [2] H. Sanchez *et al.*, "A 200 MHz 2.5 V 4 W superscalar RISC microprocessor," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, vol. 39, pp. 218-219, Feb. 1996.
- [3] R. Philhower *et al.*, "Wideband wafer-scale interconnections in a wafer scale hybrid package for a 1000 MIPS highly pipelined GaAs/AlGaAs HBT RISC," in *Proc. 1992 Int. Conf. Wafer Scale Integration*, Jan. 1992, pp. 145-154.
- [4] Y. L. LeCoz *et al.*, "Application of a floating-random-walk algorithm for extracting capacitances in a realistic HBT fast-RISC RAM cell," in *Proc. Eleventh Int. VLSI Multilevel Interconnection Conf. (VMIC)*, June 1994, pp. 542-544.
- [5] A. W. Buchwald *et al.*, "A 6-GHz integrated phase-locked loop using AlGaAs/GaAs heterojunction bipolar transistors," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1752-1761, Jan. 1992.
- [6] H. J. Greub, *Apparatus for Skew Compensating Signals*, patent 483 369 5, issued May 23, 1989.
- [7] B. Gilbert, "A new high-performance monolithic multiplier using active feedback," *IEEE J. Solid-State Circuits*, vol. SC-9, pp. 364-373, Dec. 1974.
- [8] L. Schmidt *et al.*, "New high-speed bipolar XOR gate with absolutely symmetrical circuit configuration," *Electron. Lett.*, vol. 26, pp. 430-431, 1990.
- [9] K.-S. Nah, "An adaptive clock deskew scheme and a 500 ps 32 bit register file for a high speed digital system," Ph.D. dissertation, Rensselaer Polytechnic Institute, Troy, NY, Aug. 1994.